1988
research in
computer science

A REPORT OF ACTIVITIES DURING CALENDAR YEAR 1987
COLLEGE OF ENGINEERING · UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN · UIUC
REPRINTED FROM THE SUMMARY OF ENGINEERING RESEARCH

computer science

C. W. GEAR, Head
252 Digital Computer Laboratory, 1304 W. Springfield Ave., Urbana, IL 61801

The Department of Computer Science is the largest and one of the oldest in the country. Founded in 1949 as a research group for the purpose of constructing a digital computer, the Digital Computer Laboratory (DCL) steadily evolved into a departmental unit that has been known as the Department of Computer Science since 1964. The department continues a long tradition as one of the world's leaders in computer research.

From the earliest days in the history of automatic computation, the department has been one of a small number of university groups to design complete systems of hardware as well as software. The ORDVAC computer, built in DCL and subsequently shipped to the Aberdeen Ballistic Research Laboratories, was the first working bit-parallel computer. ILLIAC I was completed following a similar design shortly thereafter and was in operation from 1952 until 1962. It was one of the first electronic computers built and owned entirely by an educational institution. The department was involved in the design and construction of ILLIACs II, III, and IV, and its faculty are now involved in the design of a new supercomputer, Cedar, in the UIUC Center for Supercomputing Research and Development. The number of faculty has steadily grown and in 1987 numbered 45. A growth to 53 faculty is planned to coincide with the occupancy of a new building for the department in 1990.

An extensive network of computing facilities and support equipment are both in the department and on campus for use in faculty as well as student research. Departmental facilities are coordinated by the Computing Research Laboratory. The most recent new campus facility is the CRAY X-MP/48 supercomputer operated at the NSF National Center for Supercomputing Applications.

The department currently operates the following shared equipment for computer science research: a 20-processor Encore Multimax with 48 MB of RAM, a Pyramid, an AT&T 3B15, 2 AT&T 3B20s, 2 DEC VAX 780s, 7 DEC VAX 750s, and several Imagen laser printers. In addition, department researchers have access to a wide array of single-user workstations including SUNs, AT&T 3B2s, Tektronix Smalltalk machines, HP 300s, Micro Vaxes, RT/PCs, and both Symbolics and TI Explorer LISP machines. Graphics research is supported by a General Electric Graphicon, a Chromatics CX 1536, a Tektronix 4129, and a film recorder. A second Multimax and Ametek and Intel hyercubes support multiprocessor research. All equipment is networked by means of a central network and 12 or so subnetworks, all of which use Ethernet and TCP/IP. The department is networked via a Proteon 80 MB/s fiber optic network to other facilities on campus. The department has connections to ARPAnet, Bitnet, CSnet, Telenet, Usenet, and NSFnet.

ARCHITECTURE AND HARDWARE

UNIFIELD Computer

W. J. Poppelbaum,* J. Jabusch, G. Johnston, J. Estell,
D. Ballman, L. Lindstrom, S. Pandya, J. Koszczuk,
J. Davies
University of Illinois

Before going over from serial processing to parallel processing, we have designed a new set of circuits using current mirrors and small swing switching through current control. We are also exploring the use of direct (unencoded) octal and wired tables for each digital position. This will allow on-line processing of 10^6 pixel matrices at TV-compatible frequencies!

Digital Computer Arithmetic

J. E. Robertson,* F. Angelotti
National Science Foundation, CCR 86-17326

This project focuses on structures for arithmetic which are in some sense "universal." A theory of decomposition of complex structures into a limited number of simple structures whose complexity is on the order of a binary half adder has been developed. The scope of the decomposition algebra is at present limited to structures involving addition and subtraction. The difficulty of design using the decomposition algebra appears to be linear with the number of input variables, whereas the difficulty of design for Boolean algebra grows exponentially with the number of input variables. The solutions resulting from the decomposition algebra give an indication of the internal connectivity of the simple basic structures.

ARTIFICIAL INTELLIGENCE

Automatic Programming

N. Dershowitz*
University of Illinois

Given a formal specification for a desired program, our goal is to transform it step-by-step into executable code. We have formulated a set of rules for the formation of conditions and iterative loops in a top-down fashion. Currently, we are investigating the use of the Knuth-Bendix completion procedure, and variants thereof, to synthesize both functional programs and logic programs from specifications and domain knowledge expressed as rewrite rules, using deductive and inductive reasoning.

Algorithmic Debugging

N. Dershowitz*
University of Illinois

We are investigating methodologies for automated testing and debugging of logic programs. Logic programs not only have clean syntax and well-understood semantics, but also exclude many details of a procedure's run-time behavior which usually complicate analysis of their logical properties. Our approach is characterized by combining symbolic execution and automatic bug locating and correction. By using logic as both a specification and a programming language, we are building a system that, when given a program along with its specifications, has the capability to analyze and reason about the program, generate test cases, isolate bugs, and perhaps suggest corrections.

Rewrite Systems

N. Dershowitz,* G. Sivakumar
National Science Foundation, DCR 85-13417

Term-rewriting systems are a general programming language with simple, elegant syntax and semantics. We have invented various useful classes of simplification orderings for proving termination and are also investigating methods of proving confluence and partial correctness. In addition, applications of rewrite systems to theorem-proving and computation are being investigated. We are participating in the implementation of the RRL rewrite-rule laboratory, for use in program verification and algebraic program specification. Conditional rewriting is an active area of current research; we are looking at its application to constraint satisfaction, as well as theoretical issues.

Advanced Inference Engine Design

K. D. Forbus*
National Science Foundation, IRI 86-57347 PYI

In collaboration with Johan de Kleer at Xerox PARC, we have developed a new inference engine based around his assumption-based truth-maintenance system that provides nearly two orders of magnitude speed-up in qualitative simulation algorithms. We are also beginning to explore the use of assumption-based truth-maintenance systems in other kinds of reasoning, such as solving textbook physics problems.

Qualitative Models for Space System Engineering

K. D. Forbus,* J. Collins, G. Skorstad
National Aeronautics and Space Administration, NAG-9137

This project investigates the potential application of qualitative physics to NASA engineering problems, including monitoring, diagnosis, and design of space station subsystems and procedures. We are developing representations for operational criteria, fault models, and continuous signals. We are also developing a prototype collection of qualitative models for fluid and thermal systems commonly found in space station subsystems.

Qualitative Reasoning About the Physical World

K. D. Forbus,* P. Nielsen, D. Decoste
U.S. Office of Naval Research, N00014-85-K-0225

This project investigates the nature of commonsense reasoning about the physical world, especially concerning physical processes such as liquid flow, motion, and boiling. Our work on qualitative process theory has resulted in techniques for representing and reasoning about a wide range of physical phenomena. We are extending the class of inferences that may be drawn with qualitative process theory, including new techniques for prediction, interpreting measurements taken across time, and the construction and analysis of plans involving the physical world. We are also developing rich qualitative models of several domains, including fluids and mechanisms.

Probabilistic Inference

A. M. Frisch*
National Aeronautics and Space Administration, NAG 1-613

Many approaches to automated reasoning with incomplete or uncertain information are either ad hoc, theoret-

---

* Denotes principal investigator.

ically unfounded, or based on unreasonable assumptions. This research aims to unify probability theory and first-order logic in a semantically well-defined formalism that will be used to analyze and design systems for reasoning under uncertainty. We will focus on developing an inference system that proceeds by computing increasingly narrow probability intervals that converge on the desired interval. Such a system can be stopped at any time to yield correct, though partial, information, thus allowing a tradeoff between the time spent on inference and the specificity of the result.

Inference with Restricted Quantification

A. M. Frisch*

University of Illinois

This research is systematically extending automated reasoning systems with the ability to reason with statements containing restricted quantification (RQ). Such statements make assertions about individuals in a certain category. Results obtained by several researchers dramatically demonstrate the enhanced efficiency of systems employing RQ. This research is developing and studying a computational technique for reasoning with RQ known as sorted unification. This project is distinguished by its emphasis on developing a general understanding of sorted unification that explains how the technique can be used in a wide range of applications, which to date include knowledge retrieval, logic programming, and parsing.

Knowledge Retrieval as Specialized Inference

A. M. Frisch*

University of Illinois

Artificial intelligence reasoning systems commonly contain a large corpus of declarative knowledge and provide facilities with which the system's components can retrieve this knowledge. This research project is studying the very nature of retrieval. Formal specifications that capture certain informal intuitions about retrieval are being developed, studied, and implemented by retrieval algorithms. This formalization is made possible by viewing a retriever as a highly specialized, limited inference engine. The principal result of this work is a demonstration that common retrieval operations such as pattern-matching and taxonomic reasoning can be integrated and jointly specified as inference operations.

Analogical Processing and Experiential Learning in Physical Domains

D. G. Gentner* (Psychology), K. D. Forbus,*
B. Falkenhainer, J. Skorstad

U.S. Office of Naval Research, N00014-85-K-0559

The goal of this project is to develop a theoretical framework that will explain how people learn from experience about the physical world. The theory so far includes four different stages of models, ranging from prototypical behaviors to expert models. The account uses Forbus' qualitative process theory to describe domain models and Gentner's structure-mapping theory of analogy and similarity to describe the computations that drive a learner from one stage to the next. While theoretical work continues, we are also designing psychological experiments and computer simulations to test aspects of the theory.

Distributed Knowledge Acquisition System

M. T. Harandi,* R. Lange, B. Buchner

U.S. Army Construction Engineering Research Laboratory, DACA 88-86-D-0011-12

This project aims to build a distributed knowledge acquisition system by which domain experts can independently contribute to the gradual construction of a communal knowledge base. The knowledge base contains entities, such as subdomains, concepts and variables, as well as relations between these entities. The basic procedure to be followed is to ask experts to enter information concerning a domain-related topic and to determine how this information affects the existing knowledge structures. In the current phase of this project we are formalizing a representation for "domain theories," an abstract mechanism for domain modeling, and the use of such theories in proving, completing, or modifying user-specified knowledge segments.

Knowledge-based Programming Assistant

M. T. Harandi,* N. Dershowitz, Q. Ning, S. Renner

IBM Corp.

This project aims to develop an intelligent programmer's workbench knowledgable in program design, coding, and debugging. Such a programming assistant system would consist of a number of expert systems, a central database, and other software tools such as program analyzers and interactive rule generators. A heuristic program debugging system and a Pascal-based bug localization tool have been developed. Work on a design and coding assistant, and the deep model of program debugging is underway.

Expert System Environments

M. T. Harandi,* R. Lange

IBM Corp.

This project studies tools and techniques for building diagnostic and interpretative expert systems. Currently, a Pascal-based expert system building tool called GPSI has been developed which includes a graphically oriented knowledge acquisition unit, knowledge-base management facilities, a rule compiler, and an inference engine. This tool is being used to develop several expert systems for a variety of domains of application. Work is underway to design and construct a graphical tool for visualization of the inference process.

Data Flow Design Aid

M. T. Harandi,* S. Bhat, R. Price

IBM Corp.

This project has developed an expert system that assists users in constructing data flow designs for programs and systems. The system is knowledge-base driven, using a rich database of data structures, data flow models, data transformations, and abstract procedures. The user participates in the design process by specifying the relationships between program components and by supplying data definitions. The system can also function as a data flow management tool, providing access to data in the knowledge base. A graphics editor for data flow design and a limited natural language interface have been developed and interfaced with a prototype of the design aid system. This research is being continued with an emphasis on planning

techniques, constraint propagation, and capabilities for rapid prototyping.

Intelligent Program Analysis and Debugging

M. T. Harandi,* Q. Ning

IBM Corp.

This project aims to develop a knowledge-based program analysis tool to provide human expert level assistance to programmers in general program understanding and debugging tasks. It is characterized by its use of a knowledge base of coding and debugging heuristics represented as program plans. With the help of the encoded human expert knowledge, this tool is able to comprehend a given program without the need for a separate specification from the programmer or an execution of the program. Based on the acquired understanding, the system can then generate paraphrases about the intentions implied in the given code and diagnose "deep" bugs that are commonly related to these intentions.

Intelligent Operating System Environment

M. T. Harandi,* S. Bhansali

U.S. Air Force Office of Scientific Research, AFOSR, 49620-86-C-0136; IBM Corp.

This project aims to build a system that when given an information or semiformal specification of a problem produces the correct sequence of OS commands (or subroutine calls) to solve the problem. Initially, we are targeting for two versions of this system. One is for the domain of numerical computation, using libraries of numerical analysis subroutines. The UNIX operating system will be the target domain for the second version. In both cases, the emphasis is on identifying the logical steps in the solution of a problem. In most cases, the solution generated by the system will be directly executable.

Integrating Multiple-Knowledge Representations and Learning Capabilities in an Expert System

R. S. Michalski,* A. B. Baskin,* T. Channic, A. Burks, J. Elliott, D. Selig, J. Zhang, P. Stefanski, L. Watanabe, S. Westfall, L. Chachere, G. Greene, K. Kaufman, H. Ko

U.S. Office of Naval Research, N00014-87-K-0874

The aim of this project is to develop a "metaexpert" system, a general-purpose inference system for implementing expert systems for various specific domains. The systems will be able to represent knowledge of human experts in different ways and will also have learning capabilities. The first version of the system has been developed and used in the development of two expert systems for agriculture (diagnosing soybean diseases and assessing cutworm damage in corn) which utilize different control and rule evaluation schemes. Recently another expert system was developed, called BABY, that provides help in treatment of prenatal babies. New research concentrates on adding inductive learning capabilities and multiple control schemes to ADVISE.

Classification of the Generic Structure of Images

S. M. Omohundro*

University of Illinois

The basic input for a vision system is a two-dimensional array of intensity values obtained from a camera or other imaging system. From this pattern of intensity variation, we want to extract three-dimensional information that will allow us to recognize the identity and state of objects and identify their placement in space. Most systems extract an *ad hoc* set of features from images to use in later processing. We have several classifications of features based on the geometry of imaging and the physics of light that should allow for a more natural recognition procedure. We use modern differential geometry and bifurcation theory to classify the characteristics of six fundamental edge types and their terminations.

Computational Geometry Structures in Learning and Recognition

S. M. Omohundro*

University of Illinois

An object recognition problem can be formulated in purely geometric terms amenable to the tools of computational geometry. An example is the need to recognize the orientation of a known polyhedron from a perspective or orthogonal projection of it to a two-dimensional plane. We have found a good solution to the one-dimensional version of this problem and it suggests two-dimensional solutions. Another situation where geometric data structures arise is in pattern recognition, where one wants to decompose a feature space into distinct classification regions. We are currently working on optical character recognition based on Voronoi diagrams. We have a new approach to learning nonlinear mappings based on computational geometry structures.

Parallel Algorithms and Architectures for Vision

S. M. Omohundro*

University of Illinois

Parallel computing is essential for fast vision. Low-level vision is naturally parallel because there are a large number of pixels in an image. High-level vision is parallel because there are a large number of features in a scene and a large number of objects in a decent visual database. There are a variety of natural parallel architectures and algorithms to parallelize low-level vision, but high-level vision requires more sophisticated ideas. We have done extensive work on the Connection Machine massively parallel computer, which is well suited for parallel symbolic computation. We are currently studying appropriate vision algorithms and architectures for this type of machine and others.

Neural Networks in Vision and Visual Cortex Simulation

S. M. Omohundro,* B. Mel

University of Illinois

We are implementing neural networks for several visual processes. We are working on a visual cortex simulator which learns to extract edge features at different orientations. These simulations should shed light on the biological mechanism for this task and suggest efficient and robust approaches for a parallel implementation. We are particularly interested in the "emergent" characteristics of the network such as the size and shape of the receptive fields. We are also comparing the efficiency, performance, and ability to learn of neural networks with that of more traditional algorithms in several vision tasks.

Analysis and Design of Error Diffusion Dithering Algorithms

S. M. Omohundro,* N. Packard,* S. Wolfram* (Physics)

University of Illinois

The problem of optimally quantizing a smooth field according to a criterion arises in many important situations. There are a large number of image display devices, such as laser printers, ink jet printers, and liquid crystal displays, which most naturally produce binary images. There are a variety of techniques known as dithering algorithms for rendering gray scale images on such devices. A good technique is based on diffusing the error that arises from quantizing one region of an image into other regions where this error is compensated for. This produces fairly good images but contains disturbing patterns at certain gray values. We are analyzing this class of algorithms using dynamical systems techniques in the hope of producing an optimal algorithm.

AI Applications to Biomedical Signal Processing

S. Ray,* J. Bentrup, D. Anhalt, D. Adams, R. Kufrin, W. Mayse, M. Mitchell, G. Schottland

Methodist Medical Center of Peoria

The Biosignals Intelligence Group centers its work on extracting and interpreting signals of biological/medical origin and interest. Multiple-channel signals, usually originating from sleeping human subjects, are processed to extract information-bearing features. Interpretations of numerous features from as many as 16 channels are applied to multivalued-logic decision formulas, developed from expert knowledge, in a typical expert-system approach, to calculate a meaning of the signals in human terms.

Currently, a system that interprets sleep disorders is in the clinical trial phase, and other systems are in progress. Basic research on rule generators, special signal-processing methods, connection machines, and feature extraction are being investigated in support of the overall research goal.

Functional Logic Programming

U.S. Reddy*

University of Illinois

Our research is directed at unifying functional programming languages (pure LISP and Scheme) and logic programming languages (Prolog). The objective is to develop a language that integrates both paradigms into a single framework. The methodology for such integration is provided by the operational mechanism of "narrowing," which extends the evaluation method of functional languages so that they behave like logic programs. We are involved in the design and implementation of a language called Scope which supports static scoping, Horn clause logic with negation, higher-order functions, environment objects, lazy evaluation, data abstraction, and polymorphic typing.

Variable Bias, Meta-Learning, and Constructive Induction

L. A. Rendell*

University of Illinois

Because it allows prediction of future events, inductive learning is a fundamental problem of artificial intelligence. When domains are uncertain, incremental, and badly behaved, two important problems are variable bias (the dynamic choice of inductive constraints) and constructive induction (the formation of new knowledge structures). Our approach to dynamically variable bias is to base it on experience with past problems. This involves extensions of similarity-based learning or meta-learning in higher-order spaces.

Like the variable-bias system, our approach to constructive induction also manages uncertainty and dynamic environments. Current work involves second-order clustering, which can tractably discover relationships.

Reasoning-based Machine Learning

L. S. Rendell*

University of Illinois

Humans can learn concepts from few examples, in diverse domains, even with incomplete information. When using these capabilities, a person reasons why an example is a member of a concept. The goal of this project is to model and implement reasoning-based learning. This could extend current limits of both similarity-based learning (SBL) and explanation-based learning (EBL). Reasoning-based learning (RBL) extends EBL since RBL will manage uncertain environments and does not need prior knowledge of the concept. The approach also extends SBL since RBL uses "semantic similarity" based on discovery of an underlying principle.

Multiple Objectives and Meta-Optimization

L. A. Rendell,* S. C. Y. Lu*

University of Illinois

When used in search problems to find an optimal evaluation function, the probabilistic learning system outperforms other techniques of optimization. The reasons for this involve multiple representations (e.g., conceptual clusters and parameterized models) and combined techniques (e.g., probabilistic induction and curve fitting).

When these techniques are extended to the problem of optimizing many goals at once, the problem becomes one of "meta-learning." One aspect of this meta-problem is to dynamically choose the best methods based on the behavior of the current primary problem, which is observed and exploited.

Automating Knowledge Acquisition Using Apprenticeship Learning Techniques

D. C. Wilkins*

U.S. Office of Naval Research, N00014-88-K-0124

A difficult well-known problem in the construction of knowledge-based expert systems is automating end-game knowledge acquisition. This refers to the problem of refining and debugging the domain-specific knowledge that is the source of power of expert systems. The most effective method that human problem solvers use to improve their expertise in knowledge-intensive domains such as medicine and engineering is an apprenticeship. The objective of this research is to develop techniques that allow an expert system to be automatically improved in an apprenticeship setting. The hypothesis being tested is that the source of power of apprenticeship learning is the ability to complete partial explanations and use to underlying domain theories.

DESIGN AUTOMATION AND GRAPHICS

Automatic Generation of VLSI Design Systems

L. G. Jones*

University of Illinois

Truly useful systems for the design of VLSI circuits should have knowledge of the target technology. For example, the system should know the specifics of design rules, signal properties, and clocking disciplines. Unfortunately, technologies are subject to change, and if a design system is tied too closely to one particular technology, it is doomed to obsolescence. The goal of this project is to provide a formal methodology for the specification of both the structure and static semantics of VLSI technologies and to develop techniques for automatically generating design systems based on these formal specifications. One of the methods under investigation is an adaptation of the attribute grammar technique used in the automatic generation of programming language systems.

Automated Design Methods for VLSI Circuits

W. J. Kubitz,* S. T. Healey

University of Illinois

Complexity equates to interconnections in ICs. In our approach, the complete chip connectivity is designed on an abstract, hierarchical, virtual multigrid. Once the best global arrangement is achieved, the chip is instantiated automatically from the abstract grid. Chips consist of modules (with or without active devices) and modules consist of cells. Modules connect via routing channels and cells connect by abutment. Modules and cells can be tiled or generated. Modules are fixed shape or variable shape. Random logic modules are generated to shape and I/O specifications. Placement deforms variable shape modules and packs the chip. Instantiation of all or part is allowed for verification purposes. Complete definition of connectivity allows early timing estimation. The approach manages

The results of research in synthetic image generation are shown here by Gregory Rogers, a graduate research assistant.

interconnections, the major source of complexity in high-density chips.

Object-oriented Graphics Environments

W. J. Kubitz,* R. H. Campbell,* G. Rogers, K. Sung, Y. Pan, J. Barrett

University of Illinois

While graphics has the potential to revolutionize the human-computer interface, it has thus far failed to do so. Impediments have been lack of sufficient quality and speed, but these are no longer limits. The limits are now primarily the primitive level at which graphics must be used and the inertia in exploiting the current high-performance systems. This project concentrates on exploiting graphics for a distributed Unix user environment by providing a graphics programming system at a high level of abstraction using the object-oriented paradigm. Areas being investigated are high-level graphics programming using graphics, graphical user interfaces, graphical user interface management systems, and design applications in an object-oriented graphics environment.

Computer-aided Design of VLSI Circuits

C. L. Liu,* J. A. Abraham, J. Cong, X. Yao

Corporate Research and Development Center, General Electric Co.

Various aspects of computer-aided design of integrated circuits are studied. Algorithms for floorplan design, compaction, global routing, and channel routing are being developed. Software packages implementing these algorithms are being prepared.

Global Routing of VLSI Circuits

C. L. Liu,* N. Hasan

Digital Equipment Corp.

Global routing algorithms for VLSI circuit design are being studied. In particular, a parallel global routing algorithm that takes into consideration terminals to be connected later is being designed. With this look-ahead capability, the algorithm performs quite well in comparison to other known algorithms.

This photo of a mirrored ball in front of a glossy block illustrates advanced computer graphics techniques. A distributed ray-tracing algorithm was used to simulate the fuzzy reflections and the soft shadows. (Photo courtesy of the department)

Algorithms and Design Tools for VLSI Design

S. Muroga,* M. E. Baddeley, K.-C. Chen, J.-M. Hsu,
J. C. Limqueco, L.-P. Lin, Z. Xiang

Reduction of design time is an urgent problem in VLSI chip design. PLAs and ROMs are extensively used as solutions in custom-designed chips. We developed a novel approach to reduce the size of a ROM. We developed PMIN, a program for absolute minimization of a standard PLA, which is faster by an order of magnitude than previous program AMIN. Also, we are developing new algorithms to reduce significantly the size of PLAs.

Design of compact MOS networks is another solution. We developed a new version of algorithm DIMN for this purpose, improving its computational efficiency by one or two orders of magnitude.

DISTRIBUTED SOFTWARE AND NETWORKS

Studies in Distributed Database Systems

G. G. Belford*
University of Illinois

Distributed databases can provide enhanced data availability to support applications such as office automation and computer-aided design and manufacturing. But first further work is needed on problems such as maintaining the consistency of replicated data in an efficient, reliable manner and supporting the continuing functioning of the system in the face of network partitioning (e.g., allowing a user to keep working on one data copy at a workstation disconnected from the network). We are attacking these and related problems by the development of provably correct, reliable protocols and by the analysis and simulation of competing protocols and algorithms. We have also designed and are implementing an experimental distributed database management system.

Design and Implementation of a Data Traffic Manager System

G. G. Belford,* J. W.-S. Liu,* E. Muroga, J. Ng, S. Shi
U.S. Army Construction Engineering Research Laboratory, DACA 88-86-D-0006-16

This project concerns the design and implementation of a prototype data traffic manager system (DTMS) which would serve as a control point to integrate several database systems owned by the U.S. Army Corps of Engineers. Typical functions supported by the prototype DTMS system include data and file transfer between database systems, retrieval of information stored in several database systems, automatic generation of reports in response to free-format user queries, and the maintenance of the mutual consistency of replicated data items. An operational version of DTMS will be designed based on our experience in the implementation of the prototype. Basic issues in the design of such a system will be investigated.

Metropolitan Networks

R. H. Campbell,* D. Reed,* M. Faiman,* G. Murakami,
L-Y. Lu, V. Russo, R. Rajagopalan
AT&T

Future optic fiber networks will interconnect tens of thousands of computers into a metropolitan-sized information sharing system. The network will support multiple communication paths transmitting many hundred megabits a second. Such a network requires advances in the hardware interface and protocols between the computers and the networks, changes in the design of operating systems, and new user and application software. This project investigates the design of 1 gigabit interfaces between computers and optic fiber networks, the performance characteristics of such a large network and its workload, and the issues which alter the design of operating systems for computers used in such systems. Experimental evaluation of certain of the properties of such a network is done with the help of Link, a distributed Unix system, and an Ametek hypercube.

Illinois Computing Laboratory for Aerospace Systems and Software (ICLASS)

R. K. Iyer,* J. W.-S. Liu,* J. A. Abraham, P. Banerjee,
G. G. Belford, R. H. Campbell, E. S. Davidson,
T. A. W. Dwyer III, W. K. Fuchs, A. Frisch, C. W. Gear,
S. N. Kamin, K. J. Lin, J. H. Patel, S. Omohundro,
D. A. Reed, B. Wah, S. Aslam, C. A. Block, R. Brouwer,
M. I. Chen, A. Choudhary, M. Devarakonda,
K. Fakhreddine, M. Gooley, A. Grimshaw, D. Grunwald,
A. Gupta, P. Haddawy, D. Hougen, S. Hsueh, W. Hwu,
K. Jenks, M. Lembeck, R. Llames, S. Lu, G. McNiven,
S. Nair, P. T. Ong, M. Schmitz, X. Song, V. Swarup,
K. Thearling, L. T. Young, A. Yu

(Also conducted in the Coordinated Science Laboratory)

The Illinois Computing Laboratory for Aerospace Systems and Software (ICLASS) is a NASA center for excellence in aerospace computing. Its research focus is in four areas: parallel architectures and algorithms, reliable and fault-tolerant computing, distributed and real-time systems, and software engineering and artificial intelligence. In the first two areas, problems being addressed include parallel algorithms for automated design systems, multiprocessor architectures for image understanding, parallel AI algorithms and software, dependable parallel architectures, highly reliable cost-effective architectures, and diagnosis of digital systems using artificial intelligence and probabilistic techniques. Problems in the other areas include reliable distributed database management systems, multicomputer system software, object-oriented data flow system, real-time control, real-time and parallel applications for the Choices system, real-time scheduling in distributed systems, abstract and executable specification, measurement-based evaluation of LISP architectures, distributed debugging, deduction in sorted logic, convergent deduction for probabilistic logic, and visual texture discrimination.

Bus-based Topologies

L. V. Kale*
University of Illinois

We are investigating use of a bus as a building block for large-scale interconnection topologies. The potential of such topologies remains largely unexplored, although they clearly present attractive features and a design flexibility that can lead to higher performance than point-to-point topologies. The $N \times N$ grid-of-buses interconnects N^2 processors with only 2 hops; but it is not scalable. We have discovered a new family of scalable interconnection topologies, where the lengths of wires do not increase with the

number of processors in the system. We intend to develop efficient routing schemes for such topologies, investigate promising variants of the topology, and study their application for parallel processing.

Dynamic Load-balancing Strategies

L. V. Kale,* W. Shu, V. Saletore

University of Illinois

A small grain of parallelism is essential and natural in executing irregular symbolic computations. The efficiency of a parallel processing system depends on how uniformly these granules of action are distributed to processors. We have developed a dynamic load-balancing scheme that speedily distributes newly created work to the "needy" processors. We are working on improving the scheme by adding a corrective redistribution component and saturation control, i.e., not moving pieces of work around while everyone has sufficient work for himself.

The memory requirement of such a system depends critically — varying between linear and exponential — on the queuing strategy, i.e., deciding which granule to execute next. We are working on strategies that minimize memory requirement without decreasing parallelism.

Research in Real-Time Systems

K.-J. Lin,* J. W.-S. Liu*

U.S. Office of Naval Research, N00014-87-K-0827

Our current research directions in real-time systems are broad. We are interested in scheduling theory, language design, operating systems implementation, database design, fault tolerance, and artificial intelligence. We have proposed the model of imprecise computation, which sacrifices the precision of the computation result in order to meet the deadline when insufficient time is available. The model allows new approaches to problems in real-time scheduling, fault tolerance, and AI.

We are implementing a system called Concord, which provides the necessary language and system support for programs to return imprecise results. The programmer is provided with a set of new primitives that can produce and handle imprecise results.

Utilizing Imprecise Results in Real-Time Systems

K.-J. Lin,* J. W.-S. Liu,* J. Y. Chung, N. Swaminathan

U.S. Office of Naval Research, N00014-87-K-0827

This project is concerned with making effective use of imprecise results returned by prematurely terminated monotone processes in hard real-time systems to ensure timeliness of system responses and to support forward recovery. A prototype system that provides programming language primitives and system support for processes to use imprecise results is being designed and implemented. Real-time scheduling algorithms and forward error recovery mechanisms based on the availability of imprecise results are being designed and evaluated.

Optimization Problems in Data Communication and Computer Networks

C. L. Liu,* J. W.-S. Liu,* J. Cong, T. Madej, X. Shen

U.S. Office of Naval Research, N00014-86-K-0416

The goal of this research effort is to study the complexity issues, to design efficient algorithms, and to analyze their performance for many important optimization problems that arise in the area of data communication and distributed computing. Problems being addressed include satellite communication, real-time scheduling and load-balancing in distributed systems, construction of minimal standby networks, reliable and robust communication networks, and optimal multiplexing schemes.

Parallel and Distributed Algorithms

J. W.-S. Liu,* S. Omohundro,* T. Madej

Hughes Aircraft Co.

This project is concerned with the design and evaluation of parallel and distributed algorithms for signal processing, data communications, memory management, pattern recognition, and object recognition. In particular, real-time aspects of the problems are addressed.

Resilient Objects

P. Ng*

University of Illinois

The goal of this research is to investigate the different strategies with which an abstract object can be made resilient to transient failures. Executions of operations on resilient objects survive failures such as node crashes and are not subjected to long interruptions. A spectrum of strategies that provide different levels of resilience as well as different hardware costs and execution performance can be used. For example, procedures operating on the object can be checkpointed or have its execution replicated on different processors. These strategies become more practical as the cost of processors decreases and distributed systems become commonplace. By capturing the resilience property in an abstract object, a standard programming interface is used to access a resilient object regardless of the strategy used to implement resilience.

ROSE — A Reliable Operating System Environment

P. Ng,* K.-J. Lin*

University of Illinois

We are implementing a modular distributed operating system, ROSE, that provides support for building reliable applications that may have real-time constraints. Our design will be structured in layers with well-defined interfaces. In the kernel layer, we intend to provide network IPC with different degrees of reliability and the capability of multicasting. Above the kernel layer, we provide highly available memory (HAM) objects whose content is always accessible despite hardware failures. In the layer above, a resilient process (RP) abstraction is provided for user processes to survive hardware failures with minimal interruption. In all layers, we also include supports for applications with real-time constraints.

INFORMATION SYSTEMS

File Management

S. M. Kuck,* R. Kaiser, Jr.

IBM Corp.

We have implemented Flexible File System, which contains both network and relational database query languages. It contains high-level commands to implement optimization

algorithms that merge relational and network query translation techniques. Flexible File System is used to support Easybase, a universal relation interface.

ER-Easy

S. M. Kuck,* R. J. A. Lewe, M. Najork
IBM Corp.

We have implemented the prototype of ER-Easy, an entity-relationship diagram design program that runs on a Sun workstation. ER-Easy also contains a user-friendly inferencing process that aids the user in producing appropriate diagrams for attributes that are over the same domain and play distinct roles. We are currently connecting it as a front-end to Easybase and investigating other database scheme design aids.

Performance Evaluation of Database Management Systems

S. M. Kuck,* Y. Pang
IBM Corp.

We are carrying out performance evaluation of a network database management system versus a relational database management system to determine the optimal implementation strategy. We have built a prototype simulator, Easyplex, of relational and network database management systems. This work has led to results in two areas. We have developed an algorithm for optimally allocating main memory buffers during the execution of a network query. We have introduced algorithms to calculate the expected number of block accesses for randomly accessing a file with buffering. Also, several noniterative approximate formulas for estimating block accesses have been derived.

High-Level Query Languages

S. M. Kuck,* S. Pax
IBM Corp.

We have added the ability to express the transitive closure to a universal relation query language. We are currently formulating equivalent versions of tuple relational calculus and relational algebra with a recursive query capability. We plan to investigate semantic-packed operators, such as *no less than*, that would allow us to pose queries such as, "Find all professors who published no less than five papers during the 1987 calendar year." Problems include query translation and optimization of queries.

Knowledge Base for Combat Engineers

M. E. Williams,* E. At-Taras
U.S. Army Construction Engineering Research Laboratory
(Conducted in the Coordinated Science Laboratory)

The objective is to develop a large, field-operated, electronic, knowledge-based information system pertaining to the combat engineer mission for easy retrieval of information in support of general or specific information objectives. Ultimately a system composed of the following three parts should emerge: (1) the necessary core resources of reference data and training material, (2) a scheme to facilitate both the information retrieval and learning under field conditions, and (3) portable equipment for accessing the stored data and the training and/or learning facilities. The initial objective of this project is to study the requirements of such a system in order to design the structure of a knowledge base and expert information system.

Intelligent Gateway Systems

M. E. Williams*
U.S. Department of Agriculture
(Conducted in the Coordinated Science Laboratory)

This project identified and evaluated the characteristics of "intelligent gateway" and user-friendly front-end systems to on-line information retrieval systems. Recommendations were made as to how an intelligent gateway could be used or developed by the NAL for the agriculture community. Scenarios for searches on agricultural problems made through intelligent gateways were developed.

Comparative Analyses of On-line Retrieval Interfaces

M. E. Williams,* M. Kinnucan, L. Smith, D. S. Cho
National Science Foundation, IST 82-0974
(Conducted in the Coordinated Science Laboratory)

This research project included a controlled experiment for the systematic comparison of end-user use of three types of retrieval system/interfaces: command-drive Boolean (CB), prompted Boolean (PB), and associative retrieval (AR). The database searched was a subset of the MEDLARS database containing information on topics of interest to laymen. We studied the patterns of end-user interaction with each of the interfaces and the outcome in terms of both retrieval system performance and user satisfaction. Recall was higher for the associative interface. GSLIS students achieved higher recall than CS students. The associative interface was considered easiest to learn and use. The test demonstrated that if a retrieval system is truly "user-friendly," users will be satisfied regardless of the type of interface.

Database Analysis

M. E. Williams,* E. At-Taras
University of Illinois

Analyses of data in the database of databases are run routinely. Analyses included number and percentage of databases by field of science, country, and sector of the economy. Statistics have been also generated regarding numbers of records within databases according to field, country, and sector of the economy. Various correlations between data items have also been generated.

The faculty and students have the department's extensive library to aid them in their research and studies.

Adding Knowledge to Databases

M. S. Winslett*

University of Illinois

This project addresses the issues that arise when one attempts to add "intelligent" reasoning facilities to a database management system. We are interested both in exact, formal approaches to these problems and in heuristic solution methods. Recent work has explored formal methods of incorporating new information into knowledge bases, given that only incomplete information is available about the state of the world. By using mathematical logic as an underpinning, these methods were made applicable to a number of problems that arise in applications within artificial intelligence, as well as to the problem of updating databases with incomplete information.

Information Systems for Engineering Design and Manufacturing

M. S. Winslett*

University of Illinois

This project addresses issues associated with the design and construction of an efficient and easy-to-use engineering information system to support engineering design activities and manufacturing. Research goals here range from efficient representations and storage and access structures for design objects and engineering processes, to the incorporation of simple artificial intelligence techniques into the information system to automate functions that, due to the high volume of data involved or the degree of precision and attention to detail required, are burdensome to humans.

NUMERICAL ANALYSIS

Computational Methods and Software for ODEs and Their Applications

C. W. Gear,* R. D. Skeel,* C. Carter, F. Juang, B. Leimkuhler, W. Mitchell, H. Tam, S. Aslam, K. Cochran, R.-L. Guo

U.S. Department of Energy, DE-FG02-87ER25026

The evolutionary behavior of large physical systems is simulated by the solution of time-dependent differential equations. These are extremely important in understanding physical and chemical processes and in the design of complex systems such as high-speed integrated circuits. The effective solution of these equations on modern parallel computers requires new approaches such as the waveform method, which permits independent processors to work on different parts of the system. Communication between the processors must be at least equal to the information flow needed to maintain stability and accuracy in the physical problem, so it is important to determine the minimum information flow required: this will determine the underlying limits of computational speed achievable no matter how great the parallelism.

Iterative Methods for Initial Value Problems

C. W. Gear,* P. E. Saylor,* R. D. Skeel,* S. Ashby, M. Holst, S. Lee, J. Keiper

National Science Foundation, DMS 87-03226

Transient analysis of large-scale models requires the solution of nonlinear equations. Often these are solved by a Newton iteration, and the resulting linear equations are solved directly using sparse techniques. The increasing size of the models suggests that iterative methods be used. Nonlinear problems that arise in transient analysis have special properties that may aid in their solution: typically they are only mildly nonlinear; the associated linear problems have all of their eigenvalues in the right-half plane; and they have special structure. We study the use of iterative methods in the special context of transient analysis and transient analysis methods in the context of iterative methods to develop techniques that will be applicable on parallel or vector computers.

Design of Algorithms for Semiconductor Simulation

T. Kerkhoven*

University of Illinois

High-quality device simulation tools are indispensable to the introduction of new VLSI technology. Three-dimensional and time-dependent modeling become necessary as device sizes shrink, speeds increase, and submicron technologies take over. We are concerned with the design of new and faster algorithms for the simulation of both the steady state and the transient behavior of silicon and gallium-arsenide (GaAs) devices. The latter part of our research is done in collaboration with K. Hess of the Coordinated Science Laboratory. Thus we examine coupled and decoupled approaches to semiconductor simulation and diverse acceleration techniques. Both theoretical and computational projects are part of this research.

A Taxonomy for Conjugate Gradient Methods

P. E. Saylor,* S. Ashby, T. Manteuffel

Los Alamos National Laboratory; U.S. Department of Energy, DE-FG02-87ER25026

A study has been made on classifying all possible conjugate gradient (CG) methods. It has resulted in a number of new CG methods including one suitable for complex symmetric matrices, important in electromagnetics and signal processing. Also, new algorithms and software have been developed for polynomial preconditioning to use with the CG method. This has been applied to the matrices arising from the simulation of groundwater flow in three dimensions in a set of numerical experiments.

Leapfrog Variants of Iterative Methods for Linear Algebraic Equations

P. E. Saylor*

University of Illinois

For two iterative methods, Richardson's method and a general second-order method, variants are derived for which only even numbered iterates are computed. The variant is called a leapfrog method, and has some advantages over certain architectures. In the case of Richardson's method, it is possible to express the final iterate in terms of only the initial approximation, a variant of the iteration called the grand leap. A set of parameters is required, and an algorithm has been derived to compute these parameters, related to algorithms to compute the weights and abscissas for Gaussian quadrature.

PARALLEL COMPUTING

Parallel Execution of Logic Programs

L. V. Kale,* U. S. Reddy,* W. Shu, B. Ramkumar
National Science Foundation, CCR 87-00988

We are studying ways of executing combinatorially explosive symbolic computations expressed as logic programs speedily and efficiently using massively parallel machines. We have developed a process model that extracts the maximal parallelism from given programs while keeping the subcomputations independent. It is based on the pursuit of both the AND and OR sources of parallelism in concert. The challenge is to reduce the overhead and produce an efficient implementation. Our current effort is focused on constructing parallel implementations of the model on various parallel machines, developing static analysis techniques to extract parallelism in logic programs, and compiling them for efficient execution.

Simulator for Parallel Computer Systems

L. V. Kale,* W. Shu, W. Saletore
University of Illinois

A simulation system is essential for empirical research on parallel computation. We are developing ORACLE, a system that simulates symbolic computations on message-passing parallel processors. Written in SIMSCRIPT and C, ORACLE models each processor and communication channel in the system. It accepts pure Prolog programs and many architectural parameters as input. It has several built-in load-balancing modules and can accept different interconnection topologies. It produces a variety of performance data such as time to completion and average processor uses. The future work includes a detailed and accurate timing module, more sophisticated buffer queuing management, memory management, and detailed communication protocol implementation.

Hypercube Operating System Design

D. A. Reed,* D. C. Grunwald, D. Bradley, D. Rudolph, B. Nazief
National Science Foundation, CCR 86-57696 PYI

In addition to features shared with uniprocessor operating systems, a hypercube operating system should also provide support for scheduling parallel tasks and internode synchronization and communication. These areas are united by their need for resource status information.

The types and amounts of status information that can readily be acquired by the nodes of a hypercube are limited. Most notably, the absence of global memory in a hypercube means that no single node can have a complete or current description of the global network state.

Based on these observations, we have designed and are testing a new hypercube operating system, "Picasso," that will provide efficient communication primitives, adaptive routing of messages to minimize communication delays, dynamic assignment of tasks to idle processors, and recovery from processor failures.

Performance Analysis and Algorithms

D. A. Reed,* D. C. Grunwald, D. Bradley, D. Rudolph, B. Nazief
National Science Foundation, CCR 86-57696 PYI

It has been a truism in computer systems design that "balanced" systems achieve the best performance. Message-passing parallel processors are no different; optimizing the performance of algorithms requires a judicious combination of node computation speed, message transmission latency, and operating system software. To quantify the balance of a hypercube design, we have developed an experimental methodology and applied the associated suite of benchmarks to most existing hypercubes.

The benchmark suite includes synthetic communication models that reflect a variety of program schemes, including CSP and functionals.

PROGRAMMING LANGUAGES, COMPILERS, AND OPERATING SYSTEMS

Real-Time Distributed Operating Systems

R. H. Campbell,* G. Johnston, K. Kenny
National Aeronautics and Space Administration, NSG-1471

An implementation of "remote objects" in Path Pascal permits portable distributed system software. A remote object specifies the synchronization for, access to, and transformations on encapsulated data that are located on a computer in a communications network. Remote procedure calls support the invocation of operations on remote objects. Distributed Path Pascal compiles to VAX 780 and 68000 machine code and has a C and UNIX binding.

An axiomatic definition of open path expressions permits verification of Path Pascal programs. Atomic actions supporting fault-tolerant operations on encapsulated data are being used in software reliability studies.

Improving Software Reuse with Object-oriented Programming

R. E. Johnson,* B. Foote
University of Illinois

The best way to reduce the cost of software is to avoid writing it. Object-oriented programming provides several ways to reuse software, including inheritance and composition. We are investigating design rules for class hierarchies and object decomposition. The best way to understand a programming paradigm is to practice it. We have built a number of object-oriented systems, including ones for laboratory control, music composition, visual programming, and text editing. A major subproject is a toolkit for building iterative graphical user interfaces.

Optimization for Object-oriented Programming Languages

R. E. Johnson,* J. Graver, L. Zurawski
University of Illinois

Object-oriented programming languages tend to be very inefficient, in large part because binding of procedure names to procedures must occur at run time instead of at compile time. We have devised a data-type system for smalltalk-like languages and are using it to build an optimizing compiler for Smalltalk. Our goal is to make Smalltalk ten times faster. We are examining a wide range of code optimization and program analysis techniques and studying the problem of debugging optimized programs.

Environments for Programming in the Many and Programming in the Large

S. M. Kaplan*

AT&T ISEP Project

Programming environments research up to now has focused mainly on programing in the small — programming support for the development of small programs. We are investigating programming in the large and programming in the many, the latter being a way to achieve the former by cooperation among programmers. We are developing tools for supporting programming in the many, and new theories and algorithms for use in these tools. We are also looking at support for other aspects of programming in the large, such as smart recompilation and software scavenging.

Bus Connection Networks

M. D. Mickunas*

University of Illinois

We are investigating the properties of various bus connection networks (viewed as generalized incidence systems, such as projective planes). The mathematical properties of such networks suggest some intriguing new techniques for performing both computation and communication. Our investigation centers on four topics: (1) discovery of simple, expandable, and inexpensive interconnection strategies which can be applied to very large networks, having possible thousands of nodes; (2) specification of simple, adaptable message routing schemes which can be implemented on those networks; (3) network survivability under partial degradation; and (4) discovery of practical applications of such networks.

Error Recovery in Compilers

M. D. Mickunas*

University of Illinois

We have devised a syntactic error recovery scheme for programming languages. The scheme is based on LR parsing and is driven by information that is directly and automatically obtainable from the LR tables. The method shows good promise for providing excellent error diagnostics. In continuing work, we wish to improve the efficiency of the method and to relate the mechanism to other parsing schemes.

Tools for Compiler Construction

M. D. Mickunas*

University of Illinois

In this project we are developing a generalized shift-reduce parser generating system which allows interactive generation and modification of noncanonical LR-style parsers. This is coupled with a tree evaluator which provides interpretive execution of an attribute-like semantics for the language. This will be used as a tool to produce parsers for unconventional languages, to experiment with error recovery techniques, and to investigate aspects of parallel parsing.

Multiprocessor Compilation

M. D. Mickunas*

University of Illinois

In this project we are investigating techniques for compiling in parallel. We have devised an extension to standard shift-reduce parsing, called "piecewise" LR parsing, which permits multiple parsers to be situated at arbitrary places in an input string. Using such piecewise LR parsers, we have demonstrated that it is not only possible to parse in a truly parallel fashion, but that significant speedup can often be achieved. Our continuing goal is to characterize the class of languages acceptable by such parsers, and to devise unconventional programming language constructs which exploit the power of the technique.

Analysis of Programs

U. S. Reddy,* C. W. Pyo

University of Illinois

Static analysis of programs provides valuable information about the execution-time behavior of programs which can be used for verification of program properties and for transformation of programs into more efficient ones. We have formulated a powerful language for expressing such analysis information based on the notion of "regular tree sets," and are currently involved in designing algorithms for automatic inference using regular tree sets. Unlike conventional program checkers (e.g., Pascal), our type of inference systems does not rely on programmer-supplied declarations for performing type checking, but generates its own information using static analysis. We have formulated an algorithm for doing type inference for functional programs and are currently working on the algorithm for logic programs.

SOFTWARE ENGINEERING

SAGA — A Project to Automate the Management of Software Production Systems

R. H. Campbell,* D. Laliberte, H. Render, R. Sum

National Aeronautics and Space Administration, NAG-1-138

Quality software is expensive and time-consuming to produce. The SAGA project is seeking improvements to the software development process by means of automation. A model lifecycle based on the PLEASE/ADA executable specification language that supports incremental refinement, validation, and verification is being prototyped in the ENCOMPASS environment. ENCOMPASS includes a test harness for executable, partially designed programs and mechanical verification aids. Advanced project management and configuration control tools are being developed as a result of these studies. These tools are based on an entity relation model of the software. A formal model of the software development process is being investigated.

Illinois Software Engineering Project

R. H. Campbell,* M. T. Harandi,* S. N. Kamin,*
J. W.-S. Liu,* R. E. Johnson,* S. M. Kaplan,* S. Bhat,
B. Cheng, I. Cheong, S. Goering, J. Graver, L-Y. Lu,
D. Hammerslag, T. Kraus, A. Silberman, W. Smith

AT&T Communications Systems

The primary objective of this cooperative project with industry is to develop a set of interrelated and integrable software tools within an open system architecture which facilitates the design, development, and maintenance of reliable, robust, and reusable software systems. Some of the tools included are language-oriented and tree editors,

revision control systems, configuration management and change control systems, project management, and engineering databases. Knowledge-based support for various software development activities including design are being developed and coordinated with the tools. Studies also include object-oriented programming methods and tools for software reuse and software executable specifications for object-oriented programs. Much of the work is being applied to C++.

EOS — A Project to Investigate the Design and Construction of Real-Time Distributed Embedded Operating Systems

R. H. Campbell,* G. Johnston, K. Kenny

National Aeronautics and Space Administration, NSG-1471

Parallel processing, high-bandwidth communications, and inexpensive hardware will have an impact on the design of future operating systems. This project investigates the design and construction of a family of real-time distributed embedded operating systems for reliable, distributed aerospace applications. The system is object-oriented and targeted to support adaptive distributed software applications on multiprocessors and networked processors. Each real-time task has a virtual memory which may be mapped onto several processors. Collections of concurrent lightweight processes in each task may invoke operations on collections of objects, some of which may be shared between tasks. Fast context switching and object-oriented design help to maximize the system's performance.

Software Design

R. H. Campbell,* N. Dershowitz,* M. T. Harandi,* R. E. Johnson,* S. M. Kaplan*

University of Illinois

The investigators are studying ways of automating and aiding the software design process using a combination of expert systems, knowledge-based approaches, formal methods, executable specifications, design heuristics, data flow analyses, and design by example.

An Environment for Program Specification, Synthesis, and Verification

M. T. Harandi,* R. H. Campbell,* L. K. Miriyala, F. Young

University of Illinois

This project involves the creation of an intelligent software design environment for construction of reliable programs. The environment will include a tool for development of formal specifications from constructive dialogues, tables, formulae, and other informal means. The specification aid will be capable of detecting missing or inconsistent information. A second aspect is the construction of tools to transform and augment specifications as a result of design decisions, and to provide assistance for implementation and code generation. These tools will check refinements for completeness and inconsistencies and provide machine assistance for avoiding common errors, generating routine code, enforcing coding standards, and formatting. Documentation will be automatically inserted in the final program.

Specification and Implementation of Software Development Environments

L. G. Jones,* S. M. Kaplan*

University of Illinois

A number of language-based environments have been implemented using attribute grammars which bind the static semantics to the syntax of the target programming language. This project examines two problems. (1) Since language specifications based on attribute grammars are often very cryptic, we are studying the specification of static semantics using higher-level formalisms, as well as the use of circular, inductive definitions and dependencies forming arbitrary graph structures to extend the expressive power of attributes. (2) As systems based on attribute grammar techniques are typically slow and require a great deal of space, improved representation and evaluation techniques are needed to reduce the inherent redundancies.

FASE3

S. N. Kamin,* T. Kraus

AT&T

We are developing a C++-based executable specification system called FASE3. Aspects of interest are: (1) FASE3 is a two-level specification language in which the abstract, or "shared," language is separated from the host-language-dependent, or "interface," language. (2) Specifications are highly abstract. (3) Specifications are usually executable, even when quantifiers occur in function definitions. FASE3 employs a sophisticated evaluation mechanism based upon a novel combinator translation.

XTED

S. N. Kamin,* D. Hammerslag

AT&T

We are developing an EMACS-like editor for trees, called XTED. We feel that previous approaches to tree-editing have placed too much emphasis on security and not enough on usability. A central tenet of our approach is to give to programmers the tools to create useful tree editors, without the constraints imposed by earlier editors. The editor will provide a (partially customizable) visual front end, and a language in which to manipulate the underlying tree structure. This language will allow us to customize the editor for various applications. It is to be expected that support for applications will not be as "deep" as in special-purpose (e.g., program) editors, but this will be compensated by greater efficiency, flexibility, and a wider range of applications.

Applications of Tree Editing

S. N. Kamin,* D. Carr, D. Hammerslag

University of Illinois

We believe that a sufficiently flexible and extensible tree editor can support a much wider range of computing activities than has previously been appreciated. It can provide the basis of a highly usable, learnable, and flexible computing environment. Typical applications are program editing and outline/document processing; more advanced applications currently under development are formal program development and program transformation systems.

Tree-based Programming

S. N. Kamin,* D. Carr

Illinois Computing Laboratory for Aeronautics Systems and Software

The use of a tree editor to edit programs can provide for abstractions not available in the underlying programming language. In our opinion, the proper way to structure programs as trees is at the level of cliches, which are parameterized program fragments. Past attempts to use cliches (e.g., Programmer's Apprentice project at MIT) have been hampered by the effort to retain a primarily textual view of programs. By making the cliche structure entirely explicit, we allow for a far simpler and more efficient editor, and, more importantly, can offer more flexibility in the construction and modification of cliches. We believe we can provide, in a more or less language-independent way, various kinds of abstraction: textual, functional, data, control.

Computing Research Laboratory

W. J. Kubitz,* R. A. Aydt, J. R. Carey, A. K. Cheng, C.-L. Heddle, A. F. Irwin, P. G. Richards, T. F. Schurg, M. Schwager

University of Illinois

The Computing Research Laboratory continues to lead the way in exploring modern concepts to provide a flexible and convenient environment for computer science research. The current system is a hierarchical network of heterogeneous research networks, each suited to the needs of specific research groups. While the primary net provides common services to all researchers, each group network consists of specialized research computers and workstations. Newest additions to the research environment are a second Encore Multimax, an Intel cube-connected processor, and a number of Unix workstations. The future promises to bring in more and more specialized processors to support the diverse research needs of the department.

FOCUS Program Derivation System

U. S. Reddy*

National Aeronautics and Space Administration, NAG 1-613

This project investigates methodologies and automated tools for formal program derivation. The FOCUS system, the test bed for these methodologies, supports the paradigm of automation-based software development and maintenance. The software designer initially specifies the behavior of a system in a formal specification language. The specification language used in the FOCUS system is a higher-order functional logic language with recursive term data structures, sets, mappings, and relations. Such a specification can be executed using the FOCUS system and validated against customer requirements. The designer then systematically transforms the specification into an efficient program using the program derivation tools provided by the FOCUS system.

THEORETICAL COMPUTER SCIENCE

Enumeration of Trees

N. Dershowitz*

University of Illinois

Various enumeration problems concerning the class of ordered trees are being studied. We have found a very general formula enumerating occurrences of patterns among these trees. Combinatorial and statistical properties of this class of trees and their application to the analysis of several algorithms are under investigation.

Arrangements of Jordan Curves

H. Edelsbrunner*

University of Illinois

Assume one has a circle and draws a collection of cords inside that circle, where a cord is a simple curve (Jordan curve) whose two endpoints lie on the circle. Assume that any two of the curves intersect in at most s points, for some constant s. We study topological and computational questions about such arrangements of curves. For example, consider two points in the circle. Is it always possible to draw a new cord through the two points that intersects each one of the old curves in at most s points? If it exists, how fast can such a new cord be computed?

Implementation Techniques for Geometric Algorithms

H. Edelsbrunner,* E. P. Mucke, H. Rosenberger

University of Illinois

Many of the geometric algorithms described in the literature make use of rather complicated data structures, use subroutines that run fast only in a theoretical sense, and assume away a bulk of special cases, each of which has to be treated separately. One part of this project consists of finding techniques that aid in implementing geometric algorithms. For example, we look into techniques that reduce the number of special cases coming up in geometric algorithms. The other part focuses on the production of implementations where the emphasis lies on testing those techniques as well as on identifying new difficulties and possible solutions.

The Complexity of Probing

H. Edelsbrunner,* S. S. Skiena

University of Illinois

The concept of "probing" can best be described if we imagine a robot that has a sensor at its disposal. This

One of professor Herbert Edelsbrunner's research interests is the decomposition of spatial figures into tetrahedra. The figure he holds has the property that the annulus it defines cannot be tetrahedrized without adding a new vertex inside. This phenomenon has no equivalent in two dimensions, where every polygon can be triangulated without adding new vertices.

sensor could be a moving finger, an x-ray gun, or what have you. This sensor can be used to collect information about the world. How would one go about exploring the environment if one were the robot? It turns out that in many cases the robot is not able to "explore" a given object in finite time, but in some interesting cases there are finite strategies. Finding such strategies and optimizing them is the main goal of this project.

Physical Design of VLSI Circuits
C. L. Liu,* X. Shen, R. Libeskind-Hadas
University of Illinois

Various aspects of the physical design of VLSI circuits are investigated including floorplan design, pin assignment, multilayer channel routing, and PLA minimization.

Inductive Inference
L. B. Pitt*
University of Illinois

The process of learning from examples may be viewed from the vantage point of computability theory: the learning task becomes that of eventually identifying some unknown total recursive function f, given only the values $f(0)$, $f(1)$, $f(2)$, etc. Whether or not a function can be successfully identified depends both on the particular computational model chosen, as well as the definition of successful identification. We are comparing the power of various formal models of computation for this type of inference problem. Our main interest is the power of randomized strategies, and the relationship between deterministic and parallel computational models. Interesting tradeoffs have been obtained, showing under what circumstances randomization may be substituted for parallelism, and vice-versa.

Computational Complexity of Learning
L. B. Pitt*
University of Illinois

We are presently investigating definitions which capture the intuitive notion of "learnable," yet are thoroughly formal in that the techniques of theoretical computer science can be applied to prove (or disprove) the feasibility of any given learning task. The feasibility of learning from examples has been determined for some domains, while for others the intractability of learning has been demonstrated. Relationships between learning problems and certain combinatorial optimization problems have been given. We are especially interested in how the availability of different types of information aids the learning task.

Analysis of Algorithms
E. M. Reingold,* A. Goldstein
University of Illinois

We are currently investigating a variety of problems in combinatorial algorithms. Included are algorithms for VLSI layout, matching and graphs, binary search trees, and self-organizing data structures.

JOURNALS AND BOOKS

Architecture and Hardware
KUCK, D., E. DAVIDSON, D. LAWRIE, and A. SAMEH. Parallel supercomputing today and the Cedar approach. *Experimental Parallel Computing Architectures* (Dongarra, ed.; Elsevier Science Publ.) 1–20 (1987).

YEW, P-C., N-F. TZENG, and D. H. LAWRIE. Distributing hot-spot addressing in large scale multiprocessors. *IEEE Trans. Comput.*, C-36:4, 388–395 (1987).

Artificial Intelligence
DERSHOWITZ, N. Termination of rewriting. *J. Symbolic Comput.*, 3:1-2, 69–115 (1987).

FORBUS, K. Interpreting observations of physical systems. IEEE Syst., Mgt., Cybernet., Spec. Issue on Causal and Diagnostic Reasoning, SMC-17:3, 350–359 (1987).

OMOHUNDRO, S. M. Efficient algorithms with neural network behavior. *Complex Syst.*, 1, 273–347 (1987).

RENDELL, L. A. Conceptual knowledge acquisition in search. *Computational Models of Learning* (Bolc, ed.; Springer-Verlag) 89–159 (1987).

WILKINS, D. C., W. J. CLANCEY, and B. J. BUCHANAN. Knowledge base refinement by monitoring abstract control knowledge. *Int. J. Man-Machine Studies*, 27:3, 281–294 (1987).

Design Automation and Graphics
CUTLER, R. B. and S. MUROGA. Derivation of minimal sums for completely specified functions. *IEEE Trans. Comput.*, C-36:3, 277–292 (1987).

LAI, H. C. and S. MUROGA. Logic networks with a minimum number of NOR (NAND) gates for parity functions of n variables. *IEEE Trans. Comput.*, C-36:2, 157–166 (1987).

LEONG, H. W. and C. L. LIU. Algorithms for permutation channel routing. *Integration — VLSI J.*, 5, 17–45 (1987).

MUROGA, S. Very large scale integration design. *Encyclopedia of Physical Science and Technology, vol. 14* (Meyers, ed.; Academic Press) 306–327 (1987).

WONG, D. F., H. W. LEONG, and C. L. LIU. PLA folding by simulated annealing. *IEEE J. Solid-State Circuits*, SC-22:2, 208–215 (1987).

Distributed Software and Networks
KAISER, G. E., S. M. KAPLAN, and J. MICALLEF. Multiuser: distributed language-based environments. *IEEE Software*, 4, 58–69 (1987).

REED, D. A. and C-K. KIM. Packet routing algorithms for integrated networks. *ACM Perform. Eval. Rev.*, 15:1, 7–15 (1987).

Numerical Analysis
LO, S-S., B. PHILIPPE, and A. H. SAMEH. A multiprocessor algorithm for the symmetric tridiagonal eigenvalue problem. *SIAM J. Sci. Stat. Comput.*, 8:2, s155–s165 (1987).

Parallel Computing
REED, D. A., L. M. ADAMS, and M. L. PATRICK. Stencils and problem partitionings: their influence on the performance of multiple processor systems. *IEEE Trans. Comput.*, C-36:7, 845–858 (1987).

REED, D. A. and R. M. FUJIMOTO. *Multicomputer Networks: Message Based Parallel Processing* (MIT Press, 1987).

REED, D. A. and D. C. GRUNWALD. The performance of multicomputer networks. *IEEE Comput.*, 20:6, 63–73 (1987).

REED, D. A., A. D. MALONY, and B. D. McCREDIE. Parallel discrete event simulation: a shared memory approach. Best paper award, *ACM Perform. Eval. Rev.*, 15:1, 36–38 (1987).

Programming Languages, Compilers, and Operating Systems

REDDY, U. S. Functional logic languages, part 1. *Graph Reduction, Lecture Notes* (Fasel and Keller, eds.; Springer-Verlag) 401–425 (1987).

Software Engineering

CAMPBELL, R. H., G. JOHNSTON, and V. RUSSO. Choices (class hierarchical open interface for custom embedded systems). *Operating Syst. Rev.*, 21:3, 9–17 (1987).

PUN, K. H. and G. G. BELFORD. Performance study of two phase locking in single-site database systems. *IEEE Trans. Software Engr.*, SE-13:12, 1311–1328 (1987).

Theoretical Computer Science

CHAZELLE, B. and H. EDELSBRUNNER. Linear space data structures for two types of range search. *Discrete Comput. Geom.*, 2, 113–126 (1987).

EDELSBRUNNER, H. *Algorithms in Combinatorial Geometry.* (Brauer, Rosenberg, and Salomaa, eds.; Springer-Verlag, 1987).

EDELSBRUNNER, H. The computational geometry column. *Bull. EATCS*, 33, 23–24 (1987).

EDELSBRUNNER, H. The computational geometry column. *Bull. EATCS*, 31, 111–114 (1987).

EDELSBRUNNER, H. Geometrics and algorithmics — a tutorial in computational geometry. *Bull. EATCS*, 32, 118–142 (1987).

EDELSBRUNNER, H. and B. CHAZELLE. An improved algorithm for constructing k-th order Voronoi diagrams. *IEEE Trans. Comput.*, C-36, 1349–1354 (1987).

EDELSBRUNNER, H. and D. P. DOBKIN. Space searching for intersecting objects. *J. Algorith.*, 8, 348–361 (1987).

EDELSBRUNNER, H. and M. H. OVERMARS. Zooming by repeated range detection. *Inform. Processing Lett.*, 24:6, 413–417 (1987).

LEWANDOWSKI, J. L. and C. L. LIU. SS/TDMA satellite communication with k-permutation switching modes. *SIAM J. Algebraic Discrete Meth.*, 8:4, 519–534 (1987).

PITT, L. A note on extending Knuth's tree estimator to directed acyclic graphs. *Inform. Processing Lett.*, 24:3, 203–206 (1987).

PITT, L. and C. H. SMITH. Probability and plurality of aggregations of learning machines. *Lecture Notes Comput. Sci.*, 267, 1–10 (1987).

SHEN, X. J. and H. EDELSBRUNNER. A tight lower bound on the size of visibility graphs. *Inform. Processing Lett.*, 26:2, 61–64 (1987).

WEI, W. D. and C. L. LIU. Rearrangement inequalities. *J. Math. Res. Exposit.*, 3, 505–510 (1987).

PROCEEDINGS AND PAPERS

Architecture and Hardware

ALLEN, T. R. and D. PADUA. Debugging parallel Fortran on a shared memory machine. *Proc. 1987 Int. Conf. on Parallel Processing* (Aug. 1987, St. Charles, Ill.) 721–727 (1987).

LEE, R., P. C. YEW, and D. H. LAWRIE. Data prefetching in shared memory multiprocessors. *Proc. 1987 Int. Conf. on Parallel Processing* (Aug. 1987, St. Charles, Ill.) 28–31 (1987).

REED, D. A., P. McKINLEY, and M. BARR. Performance analysis of switching strategies. *Proc. 1987 Symp. on Simulation of Comput. Networks* (Aug. 1987, Colorado Springs, Colo.) 130–141 (1987).

ROBERTSON, J. E. Error detection and correction for addition and subtraction through use of higher radix extensions of Hamming codes. *Proc. 8th Symp. on Comput. Arith.* (May 1987, Como, Italy) 226–229 (1987).

ROBERTSON, J. E. A new design for two's complement array multipliers. *Proc. 21st Ann. Asilomar Conf. on Signals, Syst., and Comput.* (Nov. 1987, Pacific Grove, Calif.).

Artificial Intelligence

BACHMAIR, L. and N. DERSHOWITZ. Inference rules for rewrite-based first-order theorem proving. *Proc. 2nd Symp. on Logic in Comput. Sci.* (Jun. 1987, Ithaca, N.Y.) 331–337 (1987).

COLLINS, J. and K. FORBUS. Reasoning about fluids via molecular collections. *Proc. 10th Int. Joint Conf. on Artif. Intell.* (Jul. 1987, Seattle, Wash.) 590–594 (1987).

FORBUS, K. The logic of occurrence. *Proc. 10th Int. Joint Conf. on Artif. Intell.* (Aug. 1987, Milan, Italy) 9–15 (1987).

FORBUS, K., P. NIELSEN, and B. FALTINGS. Qualitative kinematics: a framework. *Proc. 10th Int. Conf. on Artif. Intell.* (Aug. 1987, Milan, Italy) 30–35 (1987).

FRISCH, A. M. Inference without chaining. *Proc. 10th Int. Joint Conf. on Artif. Intell.* (Aug. 1987, Milan, Italy) 515–519 (1987).

HADDAWY, P. and A. M. FRISCH. Convergent deduction for probabilistic logic. *Proc. 3rd Wkshp. on Uncertainty in Artif. Intell.* (Jul. 1987, Seattle, Wash.) 278–286 (1987).

KROWITZ, A. L., L. A. RENDELL, and B. HOHENSEE. The state-space of memory clusters in the Hopfield network. *Proc. Int. Conf. on Neural Networks* (Jun. 1987, San Diego, Calif.).

LEE, W. D. and S. R. RAY. Probabilistic inference. *Proc. 3rd Conf. on AI Appl.* (Feb. 1987, Orlando, Fla.) 16–20 (1987).

PITT, L., M. KEARNS, M. LI, and L. G. VALIANT. Recent results on Boolean concept learning. *Proc. 4th Int. Wkshp. on Machine Learning* (Jun. 1987, Irvine, Calif.) 337–352 (1987).

RENDELL, L. A., R. M. SESHU, and D. K. TCHENG. Layered concept learning and dynamically-variable bias management. *Proc. 10th Int. Joint Conf. on Artif. Intell.* (Aug. 1987, Milan, Italy) 308–314 (1987).

RENDELL, L. A., R. M. SESHU, and D. K. TCHENG. More robust concept learning using dynamically-variable bias. *Proc. 4th Int. Wkshp. on Machine Learning* (Jun. 1987, Irvine, Calif.) 66–78 (1987).

WILKINS, D. C. Cognitive diagnosis of heuristic classification problem solving. *Proc. 3rd Int. Conf. on Artif. Intell. and Educ.* (May 1987; Pittsburgh, Pa.) 57 (1987).

WILKINS, D. C. The representation and use of meta-level control knowledge for apprenticeship learning. *Proc. Int.*

Wkshp. of Knowledge Representation and Organ. in Machine Learning (Oct. 1987, Geseke, W. Germany) 20 (1987).

WINSLETT, M. Validating generalized plans in the presence of incomplete information. *Proc. Nat. Conf. on Artif. Intell.* (Jul. 1987, Seattle, Wash.) 261–267 (1987).

Design Automation and Graphics

CONG, J., D. F. WONG, and C. L. LIU. A new approach to the three layer channel routing problem. *Proc. IEEE Int. Conf. on Comput.-Aided Des.* (Nov. 1987, Santa Clara, Calif.) 378–381 (1987).

HASAN, N. and C. L. LIU. A force-directed global router. *Proc. 1987 Stanford Conf. on Adv. Res. in VLSI* (Mar. 1987, Palo Alto, Calif.) 135–150 (1987).

KUBITZ, W. J. and S. T. HEALEY. Abstract routing and logic networks for custom module generation. *Proc. ACM/IEEE 24th Des. Automat. Conf.* (Jun./Jul. 1987, Miami Beach, Fla.) 230–236 (1987).

KUBITZ, W. K. and Y-C. HSU. A procedure for chip floorplanning. *Proc. 1987 IEEE Int. Symp. on Circuits and Syst.* (May 1987, Philadelphia, Pa.) 568–571 (1987).

KUBITZ, W. K., Y-C. HSU, and Y. PAN. A path selection global router. *Proc. ACM/IEEE 24th Des. Automat. Conf.* (Jun./Jul. 1987, Miami Beach, Fla.) 641–644 (1987).

MUELLER, T. R., D. F. WONG, and C. L. LIU. An enhanced bottom-up algorithm for floorplan design. *Proc. IEEE Int. Conf. on Comput.-Aided Des.* (Nov. 1987, Santa Clara, Calif.) 524–527 (1987).

WONG, D. F. and C. L. LIU. Array optimization for VLSI synthesis. *Proc. 24th Des. Automat. Conf.* (Jun. 1987, Miami, Fla.) 537–543 (1987).

WONG, D. F. and C. L. LIU. Floorplan design for rectangular and L-shaped modules. *Proc. IEEE Int. Conf. on Comput.-Aided Des.* (Nov. 1987, Santa Clara, Calif.) 520–523 (1987).

Distributed Software and Networks

GRIMSHAW, A. and J. W.-S. LIU. Mentat: an object-oriented macro data flow system. *Proc. ACM 3rd Conf. on Object-Oriented Prog. Syst., Lang., and Archit.* (Oct. 1987, Orlando, Fla.) 35–47 (1987).

GRIMSHAW, A. and J. W.-S. LIU. Mentat: a prototype macro data flow system. 2nd Wkshp. on Large Grain Parallelism (Oct. 1987, Hidden Valley, Pa.).

KAISER, G. E. and S. M. KAPLAN. Reliability in distributed programming environments. *Proc. 6th IEEE Int. Symp. on Reliability in Distrib. Syst.* (Mar. 1987, Harrisburg, Va.) 45–55 (1987).

LIN, K. J., S. NATARAJAN, and J. W.-S. LIU. Imprecise results: utilizing partial computations in real-time systems. *Proc. 8th Real-Time Syst. Symp.* (Dec. 1987, San Jose, Calif.) 210–217 (1987).

LIN, K. J., S. NATARAJAN, J. W.-S. LIU, and T. KRAUSHOFT. Concord: a system of imprecise computations. *Proc. 1987 IEEE COMPSAC* (Oct. 1987, Tokyo, Japan) 78–81 (1987).

LIU, J. W.-S., K. J. LIN, and C. L. LIU. Concord prototype system and real-time scheduling. *Proc. 4th Wkshp. on Real-Time Operating Syst.* (Jul. 1987, Cambridge, Mass.) 27–30 (1987).

LIU, J. W.-S., K. J. LIN, and S. NATARAJAN. Scheduling real-time, periodic jobs using imprecise results. *Proc. 8th*

Real-Time Syst. Symp. (Dec. 1987, San Jose, Calif.) 252–260 (1987).

LU, C., A. CHEN, and J. LIU. Reliable process migration protocols. *Proc. IEEE Infocom 87* (Mar./Apr. 1987, San Francisco, Calif.) 804–813 (1987).

MALONY, A. D., D. A. REED, and P. McGUIRE. A message passing library for shared memory parallel processors. *Proc. 1987 Int. Conf. on Parallel Processing* (Aug. 1987, St. Charles, Ill.) 739–741 (1987).

PURTILO, J., D. A. REED, and D. C. GRUNWALD. Environments for prototyping parallel algorithms. *Proc. 1987 Int. Conf. on Parallel Processing* (Aug. 1987, St. Charles, Ill.) 431–438 (1987).

REED, D. A. and C-K. KIM. Packet routing algorithms for integrated networks. 1987 ACM Sigmetrics Conf. on Meas. and Modeling of Comput. Syst. (May 1987, Banff, Alta., Canada).

Information Systems

WIEDERHOLD, G., S. CHAUDHURI, W. HASAN, M. G. WALKER, and M. WINSLETT. Architectural concepts for large knowledge bases. 1st German Conf. on Artif. Intell. (Sept. 1987, Lake Constance, W. Germany).

Numerical Analysis

KERKHOVEN, T. Polynomial acceleration of Gummel's method. *Proc. NASECODE V Conf.* (Dublin, Ireland) 255–263 (1987).

Parallel Computing

GEAR, C. W. *Selected Papers from the 2nd Conf. on Parallel Processing for Scientific Computing* (Gear and Voigt, eds.; Soc. for Indus. and Appl. Math.) 287 pp. (1987).

REED, D. A., A. D. MALONY, and B. D. McCREDIE. Parallel discrete event simulation: a shared memory approach. 1987 ACM Sigmetrics Conf. on Meas. and Modeling of Comput. Syst. (May 1987, Banff, Alta., Canada).

Software Engineering

BACHMAIR, L. and N. DERSHOWITZ. Completion for rewriting modulo a congruence. *Proc. 2nd Int. Conf. on Rewriting Tech.* (May 1987, Bordeaux, France) 192–203 (1987).

DERSHOWITZ, N. and Y-J. LEE. Deductive debugging. *Proc. 4th IEEE Symp. on Logic Prog.* (Aug. 1987, San Francisco, Calif.) 298–306 (1987).

HARANDI, M. T. and M. D. LUBARS. A design environment for software systems. *Proc. Conf. on Expert Syst. Technol. in the ADP Environ.* (Nov. 1987, Washington, D.C.).

KAPLAN, S. M., S. K. GOERING, and R. H. CAMPBELL. Supporting the software development process with attributed NLS graph grammars. *Proc. Int. Conf. on Graph Grammars* (Dec. 1986, Washington, D.C.).

LUBARS, M. D. and M. T. HARANDI. Knowledge-based software design using design schemas. *Proc. 9th Int. Conf. on Software Engr.* (Apr. 1987, Monterey, Calif.) 253–262 (1987).

TERWILLIGER, R. B. and R. H. CAMPBELL. Rigorous software development using PLEASE. *Ann. Nat. Commun. Forum* (Sept. 1987, Chicago, Ill.) 1114–1118 (1987).

TERWILLIGER, R. B. and R. H. CAMPBELL. PLEASE: a language for incremental software development. *Proc.*

4th Int. Wkshp. on Software Specification and Des. (Apr. 1987, Monterey, Calif.) 249–256 (1987).

Theoretical Computer Science

CHAZELLE, B. M., H. EDELSBRUNNER, and L. J. GUIBAS. The complexity of cutting convex polytopes. *Proc. 19th Ann. ACM Symp. Theory Comput.* (May 1987, New York, N.Y.) 66–76 (1987).

EDELSBRUNNER, H., J. PACH, J. T. SCHWARTZ, and M. SHARIR. On the lower envelope of bivariate functions and its applications. *Proc. 28th Ann. IEEE Symp. Foundat. Comput. Sci.* (Oct. 1987, Los Angeles, Calif.) 27–37 (1987).

EDELSBRUNNER, H., G. ROTE, and E. WELZL. Testing the necklace condition for shortest tours and optimal factors in the plane. 14th Int. Coll. on Automat., Lang. and Prog. (Jul. 1987, Karlsruhe, W. Germany).

PITT, L., M. KEARNS, M. LI, and L. G. VALIANT. On the learnability of Boolean formulae. *Proc. 19th Ann. ACM Symp. on Theory of Comput.* (May 1987, New York, N.Y.) 285–295 (1987).

TECHNICAL REPORTS AND OTHER DOCUMENTS*

Architecture and Hardware

KALE, L. V. The mesh superceded? Dept. rept. UIUCDCS-R-87-1362, UIUC (1987).

KONSEK, M. B., D. A. REED, and W. WATCHARAWIT-TAYAKUL. Context switching with multiple register windows: a RISC performance study. Dept. rept. UIUCDCS-R-87-1377, UIUC (1987).

OFEK, Y. The topology, algorithms and analysis of a synchronous optical hypergraph architecture. Dept. rept. UIUCDCS-R-87-1343, UIUC (1987).

Artificial Intelligence

BENTRUP, J. A., G. J. MEHLER, and J. D. RIEDESEL. INDUCE 4: a program for incrementally learning structural descriptions from examples. Dept. file UIUCDCS-F-87-958, UIUC (1987).

DERSHOWITZ, N. and Y. LEE. Deductive debugging. Dept. file UIUCDCS-F-87-961, UIUC (1987).

FALKENHAINER, B. Towards a general-purpose belief maintenance system. Dept. rept. UIUCDCS-R-87-1329, UIUC (1987).

FALKENHAINER, B., K. D. FORBUS, and D. GENTNER. The structure-mapping engine: algorithm and examples. Dept. rept. UIUCDCS-R-87-1361, UIUC (1987).

FALTINGS, B. Qualitative place vocabularies for mechanisms in configuration space. Dept. rept. UIUCDCS-R-87-1360, UIUC (1987).

FORBUS, K. D., P. NIELSEN, and B. FALTINGS. Qualitative kinematics: a framework. Dept. rept. UIUCDCS-R-87-1352, UIUC (1987).

GENTNER, D. Analogical inference and analogical access. Dept. rept. UIUCDCS-R-87-1365, UIUC (1987).

GENTNER, D. Mechanisms of analogical learning. Dept. rept. UIUCDCS-R-87-1381, UIUC (1987).

GENTNER, D. and M. JEZIORSKI. Historical shifts in

the use of analogy in science. Dept. rept. UIUCDCS-R-87-1389, UIUC (1987).

HOGGE, J. C. TIME and TPLAN user's manual. Dept. rept. UIUCDCS-R-87-1366, UIUC (1987).

HOGGE, J. C. TPLAN: a temporal interval-based planner with novel extensions. Dept. rept. UIUCDCS-R-87-1367, UIUC (1987).

KATZ, B., T. W. FERMANIAN, and R. S. MICHALSKI. AgAssistant: an experimental expert system builder for agricultural applications. Dept. file UIUCDCS-F-87-978, UIUC (1987).

MATHEUS, C. J. and W. E. HOHENSEE. Learning in artificial neural systems. Dept. rept. UIUCDCS-R-87-1394, UIUC (1987).

MEL, B. W. A computational model for neural feature extraction. Dept. rept. UIUCDCS-R-87-1357, UIUC (1987).

MICHALSKI, R. S., A. B. BASKIN, C. UHRIK, and T. CHANNIC. The ADVISE.1 meta-expert system: the general design and a technical description. Dept. file UIUCDCS-F-87-962, UIUC (1987).

NIELSEN, P. The qualitative statics of rigid bodies. Dept. rept. UIUCDCS-R-87-1354, UIUC (1987).

OMOHUNDRO, S. M. Efficient algorithms with neural network behavior. Dept. rept. UIUCDCS-R-87-1331, UIUC (1987).

PITT, L. and R. E. REINKE. Polynomial-time solvability of clustering and conceptual clustering problems: the agglomerative-hierarchical algorithm. Dept. rept. UIUCDCS-R-87-1371, UIUC (1987).

RENDELL, L. Conceptual knowledge acquisition in search. Dept. rept. UIUCDCS-R-87-1317, UIUC (1987).

RENDELL, L. Layered concept learning and its advantages. Dept. rept. UIUCDCS-R-87-1320, UIUC (1987).

RENDELL, L. Representations and models for concept learning. Dept. rept. UIUCDCS-R-87-1324, UIUC (1987).

RENDELL, L., P. BENEDICT, and H. CHO. Concept acquisition from examples: measurement of system performance and suggestions for improved design. Dept. rept. UIUCDCS-R-87-1315, UIUC (1987).

RENDELL, L., P. BENEDICT, H. CHO, and R. SESHU. Improving the design of rule-learning systems. Dept. rept. UIUCDCS-R-87-1395, UIUC (1987).

RENDELL, L., R. SESHU, and D. TCHENG. Robust concept learning using dynamically-variable bias. Dept. rept. UIUCDCS-R-87-1318, UIUC (1987).

Design Automation and Graphics

JONES, L. G. Incremental compaction of flat symbolic IC layouts. Dept. rept. UIUCDCS-R-87-1386, UIUC (1987).

Distributed Software and Networks

NG, P. A commit protocol for resilient transaction. Dept. rept. UIUCDCS-R-87-1374, UIUC (1987).

NG, P. Using histories to implement atomic objects. Dept. rept. UIUCDCS-R-87-1392, UIUC (1987).

Information Systems

JOHN, R. and A. LEWE. ER-Easy. Dept. rept. UIUCDCS-R-87-1338, UIUC (1987).

---

* See also Theses.

Numerical Analysis

ASHBY, S. F. Polynomial preconditioning for conjugate gradient methods. Dept. rept. U1UCDCS-R-87-1355, U1UC (1987).

CHRONOPOULOS, A. T. and C. W. GEAR. Implementation of preconditioned s-step conjugate gradient methods on a multiprocessor system with memory hierarchy. Dept. rept. U1UCDCS-R-87-1347, U1UC (1987).

CHRONOPOULOS, A. T. and C. W. GEAR. Implementation of s-step methods on parallel vector architectures. Dept. rept. UIUCDCS-R-87-1346, U1UC (1987).

CHRONOPOULOS, A. T. and C. W. GEAR. S-step iterative methods for symmetric linear systems. Dept. rept. U1UCDCS-R-87-1345, UIUC (1987).

GEAR, C. W. Condensed chemical kinetic equations. Dept. rept. U1UCDCS-R-87-1383, U1UC (1987).

GEAR, C. W. DAEs: ODEs with constraints and invariants. Dept. rept. U1UCDCS-R-87-1385, U1UC (1987).

GEAR, C. W. Explicit stiff stability via splitting and the parallel solution of ODEs. Dept. rept. U1UCDCS-R-87-1328, UIUC (1987).

GEAR, C. W. Parallel methods for ordinary differential equations. Dept. rept. U1UCDCS-R-87-1369, U1UC (1987).

KERKHOVEN, T. A spectral analysis of the decoupling algorithm for semiconductor simulation. Dept. rept. U1UCDCS-R-87-1323, U1UC (1987).

KERKHOVEN, T. and Y. SAAD. Acceleration techniques for decoupling algorithms in semiconductor simulation. Dept. rept. U1UCDCS-R-87-1363, U1UC (1987).

LEIMKUHLER, B. J., L. R. PETZOLD, and C. W. GEAR. On obtaining a consistent set of initial values for a system of differential-algebraic equations. Dept. rept. U1UCDCS-R-87-1344, U1UC (1987).

MITCHELL, W. F. A comparison of adaptive refinement techniques for elliptic problems. Dept. rept. U1UCDCS-R-87-1375, U1UC (1987).

SAYLOR, P. E. Leapfrog variants of iterative methods for linear algebraic equations. Dept. rept. U1UCDCS-R-87-1373, U1UC (1987).

SKEEL, R. D. and M. BERZINS. A method for the spatial discretization of parabolic equations in one space variable. Dept. rept. UIUCDCS-R-87-1319, U1UC (1987).

SULEIMAN, M. and C. W. GEAR. Treating a single, stiff, second-order ODE directly. Dept. rept. UIUCDCS-R-87-1391, U1UC (1987).

TAM, H. The shifted Adams method for ordinary differential equations. Dept. rept. U1UCDCS-R-87-1376, U1UC (1987).

Parallel Computing

CAMPBELL, R., V. RUSSO, and G. JOHNSTON. The design of a multiprocessor operating system. Dept. rept. U1UCDCS-R-87-1388, U1UC (1987).

KALE, L. V. Comparing the performance of two dynamic load distribution methods. Dept. rept. U1UCDCS-R-87-1387, U1UC (1987).

Programming Languages, Compilers, and Operating Systems

CHUNG, J., J. W.-S. LIU, and K. LIN. Scheduling periodic jobs using imprecise results. Dept. rept. U1UCDCS-R-87-1307, U1UC (1987).

JOHNSON, R. A type system for Smalltalk: an optimizer's view. Dept. rept. U1UCDCS-R-87-1336, U1UC (1987).

KALE, L. V. "Completeness" and "full parallelism" of parallel logic programming schemes. Dept. rept. U1UCDCS-R-87-1321, U1UC (1987).

KAPLAN, S. M. Incremental attribute evaluation on node-label controlled graphs. Dept. rept. U1UCDCS-R-87-1309, U1UC (1987).

Software Engineering

BENZINGER, L. A. An abstract model for the stepwise development of programs. Dept. rept. U1UCDCS-R-87-1335, U1UC (1987).

BENZINGER, L. A. A model and a method for the stepwise development of verified programs. Dept. rept. U1UCDCS-R-87-1339, U1UC (1987).

CAMPBELL, R. H., H. RENDER, R. N. SUM, JR., and R. TERWILLIGER. Automating the software development process. Dept. rept. U1UCDCS-R-87-1333, U1UC (1987).

SUM, R. N., JR. A summary of the software development cycle of the AT&T System 75 in Middletown, N.J. Dept. rept. U1UCDCS-R-87-1332, U1UC (1987).

TERWILLIGER, R. B. Knowledge-based development in ENCOMPASS. Dept. rept. U1UCDCS-R-87-1334, U1UC (1987).

Theoretical Computer Science

BOARD, R. A. and L. PITT. Semi-supervised learning. Dept. rept. U1UCDCS-R-87-1372, U1UC (1987).

CHAZELLE, B., H. EDELSBRUNNER, and L. J. GUIBAS. The complexity of cutting complexes. Dept. rept. U1UCDCS-R-87-1325, U1UC (1987).

EDELSBRUNNER, H. Minimum diameter partitioning of a planar point set in quadratic time. Dept. rept. U1UCDCS-R-87-1378, U1UC (1987).

EDELSBRUNNER, H. The upper envelope of piecewise linear functions: tight bounds on the number of faces. Dept. rept. U1UCDCS-R-87-1396, U1UC (1987).

EDELSBRUNNER, H., L. J. GUIBAS, and M. SHARIR. The upper envelope of piecewise linear functions: algorithms and applications. Dept. rept. U1UCDCS-R-87-1390, U1UC (1987).

EDELSBRUNNER, H. and E. P. MUCKE. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. Dept. rept. U1UCDCS-R-87-1393, U1UC (1987).

EDELSBRUNNER, H., A. D. ROBISON, and X. SHEN. Covering convex sets with non-overlapping polygons. Dept. rept. U1UCDCS-R-87-1364, U1UC (1987).

EDELSBRUNNER, H. and M. SHARIR. The maximum number of ways to stab n convex non-intersecting objects in the plane is $2n$-2^1. Dept. rept. U1UCDCS-R-87-1322, U1UC (1987).

SKIENA, S. S. Probing convex polygons with half-planes. Dept. rept. U1UCDCS-R-87-1380, U1UC (1987).

THESES

Architecture and Hardware

ANGELOTTI, F. W. The design of elementary multiplier modules for use in symmetric parallel multipliers. M.S. thesis, J. E. Robertson, adviser (1987).

CHEN, C. Design and analysis of fault-tolerant processor arrays for numerical applications. Ph.D. thesis, J. A. Abraham, adviser (1987).

CIESIELSKI, S. M. The UTSCM Tellabs' university test station control monitor. M.S. thesis, M. Faiman, adviser (1987).

DOLLAS, A. Architecture and applications of a Unifield-type computer. Ph.D. thesis, W. J. Poppelbaum, adviser; also, dept. rept. UIUCDCS-R-87-1348, UIUC (1987).

ESTELL, J. An analog to N-state converter. M.S. thesis, W. J. Poppelbaum, adviser (1987).

HARRIS, J. Real-time rhythm analysis of synthesizer keyboard performances for the IMS live music editing system. M.C.S. project, M. Faiman, adviser (1987).

HOEFFER, J. Concurrent structural error detection for software. M.S. thesis, W. K. Fuchs, adviser (1987).

HSUEH, M. Measurement based reliability/performability models. Ph.D. thesis, R. K. Iyer, adviser (1987).

McGUIRE, P. J. A measurement-based study of concurrency in a multiprocessor. M.S. thesis, R. K. Iyer, adviser (1987).

MARTIN, S. R. A hardware interface to connect GPIB to multibus. M.C.S. project, S. R. Ray, adviser (1987).

O'TOOLE, C. Table look-up addition and multiplication in a PSP environment. M.S. thesis, W. J. Poppelbaum, adviser (1987).

PINO, G. A. D. Arithmetic operations for unary representation using PSP. M.S. thesis, W. J. Poppelbaum, adviser (1987).

PRITCHARD, F. L. Implementing exact real arithmetic. M.S. thesis, S. N. Kamin, adviser (1987).

VANZANDT, L. L. A multiprocessor system for sleep disorders interpretation. M.S. thesis, S. R. Ray, adviser (1987).

Artificial Intelligence

ALAGAPPAN, V. Intelligent aid for datatype specification. M.S. thesis, M. T. Harandi, adviser (1987).

CHIEN, S. A. Simplifications in temporal persistence: an approach to the intractable domain theory problem explanation-based learning. M.S. thesis, G. F. DeJong, adviser (1987).

CONNOR, J. M. Development and analysis of an expert system for concrete pavement evaluation and rehabilitation and an automated knowledge acquisition system. M.C.S. project, M. T. Harandi, adviser (1987).

GATTON, T. M. Decision models for expert system development: designing for multi-level user domain knowledge. M.C.S. project, M. T. Harandi, adviser (1987).

GVILLO, D. W. A hybrid symbolic-procedural adaptive methodology for stochastic simulations. M.S. thesis, M. Ragheb, adviser (1987).

HADDAWY, P. A variable precision logic inference system employing the Dempster-Shafer uncertainty calculus. M.S. thesis, R. S. Michalski, adviser (1987).

HOGGE, J. C. The compilation of planning operators from qualitative process theory models. M.S. thesis, K. D. Forbus, adviser; also, dept. rept. UIUCDCS-R-87-1368, UIUC (1987).

KEIFER, K. M. Computer classification of ear-oximetry oxygen desaturation studies. M.C.S. project, S. R. Ray, adviser (1987).

KROWITZ, A. L. Generalization and categorization in neural networks. M.S. thesis, L. A. Rendell, adviser (1987).

LIN, Y. An expert system for VLSI layout. Ph.D. thesis, D. D. Gajski, adviser (1987).

LIU, A. C. A variable-content textual item construction tool for computer-based instruction. M.S. thesis, M. Siegel and H. G. Friedman, advisers (1987).

LUBARS, M. A knowledge-based design aid for the construction of software systems. Ph.D. thesis, M. T. Harandi, adviser (1987).

MATHEUS, C. J. Conceptual purpose: implications for representation and learning in machines and humans. M.S. thesis, L. A. Rendell, adviser; also, dept. rept. UIUCDCS-R-87-1370, UIUC (1987).

NICHOLL, S. S. Language acquisition by a computer program based on first-order logic. M.S. thesis, R. E. Johnson, adviser (1987).

O'RORKE, P. Explanation-based learning via constraint posting and propagation. Ph.D. thesis, G. F. DeJong, adviser (1987).

PATRONIK, J. P. An intelligent information-sharing system. M.C.S. project, M. Waugh and K. D. Forbus, advisers (1987).

POLLACK, J. B. On connectionist models of natural language processing. Ph.D. thesis, D. Waltz, adviser (1987).

RAVLIN, S. A computer model of effective reactions to goal-relevant events. M.C.S. project, A. Ortony and G. F. DeJong, advisers (1987).

SMITH, B. D. A process-oriented qualitative analysis of electronics. M.S. thesis, K. D. Forbus, adviser (1987).

SONG, Y. The lexicon of a natural language processing system: INLAPS. M.S. thesis, R. E. Johnson, adviser (1987).

TALCOTT, R. D. Visualization of the inferencing process in expert systems. M.S. thesis, M. T. Harandi, adviser (1987).

WHITEHALL, B. L. Substructure discovery in executed action sequence. M.S. thesis, R. E. Stepp, adviser (1987).

Design Automation and Graphics

BADDELEY, M. E. Pimac: interactive placement and interconnection of MOS macro cells. M.S. thesis, S. Muroga, adviser (1987).

BLOSTEIN, D. Recovering the orientation of textured surfaces in natural scenes. Ph.D. thesis, N. Ahuja, adviser (1987).

CHEN, K. Program PMIN for PLA minimization. M.S. thesis, S. Muroga, adviser (1987).

CONG, J. A new approach to three layer channel routing. M.S. thesis, C. L. Liu, adviser (1987).

COOK, D. J. Orpheus: a music notation program. M.S. thesis, W. J. Kubitz, adviser (1987).

DUNN, W. M. Imagedit: a raster image editor. M.S. thesis, D. D. Hearn, adviser (1987).

GAUSS, S. D. Computer animation of human figures. M.S. thesis, D. D. Hearn, adviser (1987).

GRANSTON, E. D. Sled: a multi-color, text-and-graphics slide editor. M.S. thesis, W. J. Kubitz, adviser (1987).

HASAN, N. A force-directed global router. M.S. thesis, C. L. Liu, adviser (1987).

HEALEY, S. T. Abstract partitioning and routing of logic networks for custom module generation. Ph.D. thesis, W. J. Kubitz, adviser; also, dept. rept. UIUCDCS-R-87-1326, UIUC (1987).

HILS, D. D. A graphic interface to Pascal. M.S. thesis, D. D. Hearn, adviser (1987).

HOFF, W. Surfaces from stereo: an integrated approach. Ph.D. thesis, N. Ahuja, adviser (1987).

HSU, Y. Floor planning and global routing in an automated chip and design system. Ph.D. thesis, W. J. Kubitz, adviser; also, dept. rept. UIUCDCS-R-87-1353, UIUC (1987).

KIMBALL, M. P. The Oolong system — interface facilities. M.S. thesis, C. L. Liu, adviser (1987).

LAI, F. Rule based circuit optimization for CMOS VLSI. Ph.D. thesis, T. N. Trick, adviser (1987).

MUELLER, T. A fast algorithm for design or compaction of floorplans. M.S. thesis, C. L. Liu, adviser (1987).

POHLENZ, P. M. A CMOS gate matrix cell generator. M.S. thesis, W. J. Kubitz, adviser (1987).

PRASITJUTRAKUL, S. A CMOS complex cell layout generator with constraint using rule-based approach. M.S. thesis, D. D. Gajski, adviser (1987).

RECKWERDT, W. A. Dynamic abutement of partitions for automated VLSI layout generation. M.S. thesis, W. J. Kubitz, adviser (1987).

SMITH, K. J. High speed digital circuit simulator. M.S. thesis, W. J. Kubitz and H. Merkelo, advisers (1987).

WASHBURN, M. F. User interface design and POLO/finite graphics upgrade. M.C.S. project, W. J. Kubitz and L. A. Lopez, advisers (1987).

WONG, M. Algorithmic aspects of VLSI circuit layout. Ph.D. thesis, C. L. Liu, adviser; also, dept. rept. UIUCDCS-R-87-1337, UIUC (1987).

Distributed Software and Networks

BAUER, A. M. P. Development of a network and data collection system. M.S. thesis, G. G. Belford, adviser (1987).

BERNSTEIN, D. Alternative database structures for a geographic information system. M.C.S. project, G. G. Belford, adviser (1987).

BUTLER, D. Design and implementation of TCP/IP on the AT&T 3B2 computer. M.S. thesis, J. W.-S. Liu, adviser (1987).

CALLAHAN, J. On debugging real-time programs in the distributed environment. M.S. thesis, K-J. Lin, adviser (1987).

CARBALLAL, D. J. A probe-based deadlock manager for Dtroll. M.S. thesis, G. G. Belford, adviser (1987).

CHANG, J. The data dictionary module of the Dtroll system. M.S. thesis, G. G. Belford, adviser (1987).

CHEN, R. A dynamic job scheduling system for a UNIX based distributed system. M.S. thesis, W. J. Poppelbaum, adviser; also, dept. rept. UIUCDCS-R-87-1382, UIUC (1987).

CHEN, J. E. M. Three-phase commit protocol in the Dtroll system. M.S. thesis, G. G. Belford, adviser (1987).

CHUNG, J. Transparent file control system. M.S. thesis, J. W.-S. Liu, adviser (1987).

GAULKE, D. A. A distributed graphical user interface using reusable components. M.S. thesis, R. E. Johnson, adviser (1987).

JERVIS, M. Communications for the Dtroll distributed database system. M.C.S. project, G. G. Belford, adviser (1987).

KIM, C. Integrated switching networks: a performance study. Ph.D. thesis, D. A. Reed, adviser; also, dept. rept. UIUCDCS-R-87-1359, UIUC (1987).

KIM, J. L. A network surveillance mechanism to support partitioned operation in distributed database systems. Ph.D. thesis, G. G. Belford, adviser; also, dept. rept. UIUCDCS-R-87-1350, UIUC (1987).

KRAUSKOPF, T. K. TCP/IP for single-tasking microcomputers. M.S. thesis, J. W.-S. Liu, adviser (1987).

KRENSKY, A. Supporting distributed applications in Smalltalk 80. M.S. thesis, R. E. Johnson, adviser (1987).

MILLER, A. R. Nonpreemptive run-time scheduling issues on a multitasked, multiprogrammed multiprocessor with dependencies, bidimensional tasks, folding, and dynamic graphs. Ph.D. thesis, D. H. Lawrie and P. C. Yew, advisers (1987).

MILLER, G. R. Process suspension, resumption, and migration in the UNIX operating system. M.S. thesis, J. W.-S. Liu, adviser (1987).

MORALES, H. A. The lock manager in the Dtroll system. M.S. thesis, G. G. Belford, adviser (1987).

PIPPIN, P. E. Applications of expert systems to distributed system fault handling. M.S. thesis, J. W.-S. Liu, adviser (1987).

PRASTEIN, M. L. Precedence-constrained scheduling with minimum time and communication. M.S. thesis, M. C. Loui, adviser (1987).

RAFFLE, S. L. Implementation of process suspension, resumption, and migration in the UNIX 4.3BSD operating system. M.S. thesis, J. W.-S. Liu, adviser (1987).

SHULMAN, C. J. Distributed Troll and distributed Ingres: a critical comparison. M.S. thesis, G. G. Belford, adviser (1987).

Information Systems

ABELES, C. Computer application for University of Illinois residence halls mail system. M.C.S. project, G. G. Belford, adviser (1987).

FORTNER, M. Managing instruction on the University of Illinois Cluster system. M.C.S. project, W. Golden and H. G. Friedman, advisers (1987).

KAISER, R. A. FFS — support for a relational network database management system. M.S. thesis, S. M. Kuck, adviser; also, dept. rept. UIUCDCS-R-87-1358, UIUC (1987).

LEWIS, D. J. A PC accounting system for the Department of Accountancy at the University of Illinois. M.C.S. project, G. G. Belford and R. Ziegler, advisers (1987).

MANZO, W. A. Design of a physician personal workstation: implementation of a prototype laboratory window. M.S. thesis, A. B. Baskin, adviser (1987).

NYHOLT, M. P. Interactive video evaluation project. M.C.S. project, J. R. Dennis, adviser (1987).

PENICOOK, B. R. A report on computerization of the reservation system at the campus visitors center. M.C.S. project, G. G. Belford, adviser (1987).

VOORHIS, A. A. V. A database system for scheduling oral and written quizzes on laboratory experiments in

Chemistry 385. M.C.S. project, G. G. Belford and L. Belford, advisers (1987).

Numerical Analysis

CHRONOPOULOS, A. A class of parallel iterative methods implemented on multiprocessors. Ph.D. thesis, C. W. Gear, adviser (1987).

GODMAN, D. E. Simulated annealing applied to a problem in a combinatorial optimization. M.S. thesis, R. D. Skeel, adviser; also, dept. rept. UIUCDCS-R-87-1342, UIUC (1987).

SCHAEFFER, M. J. A polynomial based iterative method for linear parabolic equations. Ph.D. thesis, A. H. Sameh, adviser (1987).

Parallel Computing

BRADY, K. Solving partial differential equations: a message passing approach. M.S. thesis, D. A. Reed, adviser (1987).

CARROLL, M. J. A comparison of two dynamic load balancing schemes for multiprocessor systems. M.S. thesis, L. V. Kale, adviser (1987).

DEFEND, D. E. Parallel ray tracing in a vector-multiprocessing environment. M.S. thesis, A. H. Sameh, adviser (1987).

KWOK, A. A performance analysis of architectural scalability. Ph.D. thesis, D. A. Reed, adviser (1987).

LEE, R. L. The effectiveness of caches and data prefetch buffers in large-scale shared memory multiprocessors. Ph.D. thesis, D. H. Lawrie and P. C. Yew, advisers (1987).

Programming Languages, Compilers, Operating Systems

CHEONG, I. K. Porting link to the 4.2 BSD UNIX system. M.S. thesis, R. H. Campbell, adviser (1987).

ESSICK, R. B. The cross-architecture procedure call. Ph.D. thesis, R. H. Campbell, adviser; also, dept. rept. UIUCDCS-R-87-1340, UIUC (1987).

GRAVER, J. Adding type specification and type-checking capabilities to Smalltalk 80. M.S. thesis, R. E. Johnson, adviser (1987).

LOPEZ, F. The parallel production system. M.S. thesis, R. E. Johnson, adviser (1987).

NESS-COHN, R. L. A performance comparison of two multiprocessing UNIX system kernels. M.S. thesis, R. H. Campbell, adviser (1987).

PANGRLE, B. M. A behavioral compiler for intelligent silicon compilation. Ph.D. thesis, D. D. Gajski, adviser (1987).

POSTIGLIONE, G. Avoiding common problems with C language. M.S. thesis, G. G. Belford, adviser (1987).

RENT, C. E. Augmenting Smalltalk with Players. M.C.S. project, R. E. Johnson, adviser (1987).

ROBISON, A. A functional programming interpreter. M.S. thesis, R. H. Campbell, adviser; also, dept. rept. UIUCDCS-R-87-1327, UIUC (1987).

RUSSO, V. Link: a kernel based distributed UNIX. M.S. thesis, R. H. Campbell, adviser (1987).

TAVERNINI, V. E. Translating natural semantic specifications attribute grammars. M.S. thesis, S. M. Kaplan, adviser (1987).

WIEGAND, J. An object-oriented code optimizer and generator. M.S. thesis, R. E. Johnson, adviser; also, dept. rept. UIUCDCS-R-87-1351, UIUC (1987).

ZURAWSKI, L. Experimental results with a multi-process parsing method. M.S. thesis, M. D. Mickunas, adviser (1987).

Software Engineering

BESHERS, G. Regular right part grammars and maintained and constructor attributes in language based editors. Ph.D. thesis, R. H. Campbell, adviser (1987).

BOLDT, D. M. System's requirements definition for the C report writer. M.C.S. project, R. H. Campbell, adviser (1987).

CHENG, B. H. Seed: semantic editor. M.S. thesis, S. M. Kaplan, adviser; also, dept. rept. UIUCDCS-R-87-1341, UIUC (1987).

ERICSON, S. D. A framework for object-oriented interactive box and life environments. M.S. thesis, R. E. Johnson, adviser; also, dept. rept. UIUCDCS-R-87-1384, UIUC (1987).

McEWAN, S. D. Tools for architecture simulation and program development. M.S. thesis, D. A. Reed, adviser (1987).

McKENZIE, J. D. Prototyping in the development of software systems. M.S. thesis, D. A. Reed, adviser (1987).

TERWILLIGER, R. Encompass: an environment for incremental software development using executable, large-based specifications. Ph.D. thesis, R. H. Campbell, adviser; also, dept. rept. UIUCDCS-R-87-1356, UIUC (1987).

TOMASKO, K. S. Implementing a spreadsheet using Smalltalk 80 model/view/controller. M.S. thesis, R. E. Johnson, adviser (1987).

VADODARIA, M. V. Formal specification of database schemas and constraints using attribute grammars. M.S. thesis, G. G. Belford, adviser (1987).

WHITELEDGE, J. R. An interface for an optimizer in the highly interactive environment of Smalltalk. M.S. thesis, R. E. Johnson, adviser (1987).

Theoretical Computer Science

BACHMAIR, L. Proof methods for equational theories. Ph.D. thesis, N. Dershowitz, adviser (1987).

CARNEY, D. M. A performance evaluation of lazy heapsort. M.S. thesis, E. M. Reingold, adviser (1987).

AWARDS* AND HONORS

Geneva G. Belford

IEEE Computer Society Distinguished Visitor, 1982–85

Halliburton Engineering Education Leadership Award, College of Engineering, UIUC, 1986

Nachum Dershowitz

Professeur Associe, Ministry of National Education France, 1983 at the Universite de Paris–Sud, Orsay, France

Beckman Associate, Center for Advanced Study, UIUC, 1986–87

Xerox Award for Faculty Research, College of Engineering, UIUC, 1986

---

* Only those awarded in the past five years.

Herbert Edelsbrunner

Amoco Faculty Development Award, 1985–87
Editor, *Journal of Discrete & Computational Geometry*, 1985–
Editor, *Journal of the Association for Computing Machinery*, 1987–

Kenneth D. Forbus

Arnold O. Beckman Research Award, UIUC, 1985
IBM Faculty Development Award, 1985–86
NSF Presidential Young Investigator Award, 1987

Charles W. Gear

Fellow, Institute of Electrical and Electronics Engineers
Fellow, American Association for the Advancement of Science
Editor, *Journal of Computational and Applied Mathematics*, 1982–
Editor, *Transactions of the Society of Computer Simulation*, 1983–
Managing Editor, *SIAM Journal on Scientific and Statistical Computing*, 1985–
President, Society for Industrial and Applied Mathematics (SIAM), 1987–88
Chairman of Board, Universities Space Research Association, 1985
Doctor of Honour, Royal Institute of Technology, Sweden, 1987

Mehdi T. Harandi

Editor-in-Chief, *International Journal of Expert Systems: Research and Applications*, 1986–

Clyde P. Kruskal

IBM Faculty Development Award, 1983–85

David J. Kuck

Fellow, Institute of Electrical and Electronics Engineers
Outstanding Paper Award, International Conference on Parallel Processing, 1986
Emanuel R. Piore Award, Institute of Electrical and Electronics Engineers, 1987

Sharon M. Kuck

IBM Faculty Development Award, 1985–87

Duncan H. Lawrie

Fellow, Institute of Electrical and Electronics Engineers
Secretary, Board of Governors, Institute of Electrical and Electronics Engineers-CS, 1987

Chung L. Liu

Fellow, Institute of Electrical and Electronics Engineers
Everitt Award for Teaching Excellence, College of Engineering, UIUC, 1983

Associate, Center for Advanced Study, UIUC, 1983
Certificate of Merit, Dads Association, UIUC, 1983
University Scholar Award, UIUC, 1985
Best Paper Award, ACM/IEEE Design Automation Conference, 1986
J. S. Guggenheim Foundation Fellow, 1987

Jane W.-S. Liu

Certificate of Appreciation, IEEE Computer Society, 1983

Ryszard S. Michalski

Member, Polish Institute of Arts and Sciences
Associate, Center for Advanced Study, UIUC, 1982–83
Distinguished Visitor, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1983–85
International Chair in Informatics, University of Mons, Belgium, 1985

Saburo Muroga

Fellow, Institute of Electrical and Electronics Engineers

W. J. Poppelbaum

Fellow, Institute of Electrical and Electronics Engineers

Daniel A. Reed

IBM Faculty Development Award, 1984
NSF Presidential Young Investigator Award, 1987

Edward M. Reingold

Principal Lecturer, Fifth Biennial North Central Section MAA Summer Seminar, St. Olaf College, 1985

Ahmed H. Sameh

Editor, *IEEE Transactions on Computers*, 1983–

David C. Wilkins

Technical Editor, *AI Magazine*, 1986–

Martha E. Williams

Fellow, American Association for the Advancement of Science
Honorary Fellow, Institute of Information Scientists, U.K.
Chairman of the Board, Engineering Information, Inc.
Award of Merit, American Society for Information Science, 1984
President-elect, American Society for Information Science, 1986
Editor, *Online Review*, 1977–
Editor, *Computer Readable Databases*, 1976–
Editor, *Annual Review of Information Science and Technology*, 1975–